Network Monitoring from Your Smartphone p. 69

Windows IT Pro
A PENTON PUBLICATION
APRIL 2011 | WINDOWSITPRO.COM | WE'RE I

Sharpen Your Cloud/Mobile/ Virtualization Skills
See details p. 18

Address Cloud Security: Federated Identity Fundamentals p. 25

Exchange 2010:
Role-Based Access Control p. 30

3 Steps to BitLocker Deployment p. 36

VMware vSphere PowerCLI:
Manage VMware with PowerShell p. 39

Create a VDI Solution p. 43

Windows 7 Migration with the Microsoft Assessment and Planning Toolkit p. 48

Manage Access-Based Enumeration from the Command Line p. 51

Namespace Planning for Exchange Server 2010 p. 54

SharePoint 2010 Disaster Recovery p. 59

Netezza. Up and running in 24 hours, not 24 days.

Get set up in hours instead of days, and start counting returns in minutes instead of hours. All with IBM's Netezza data warehouse appliance for high-performance analytics. It gives you analytics reports at supersonic speeds. At a fraction of the cost of Oracle Exadata. Get real, actionable business results fast.

ibm.com/facts

COST comparison based on publicly available information as of 2/10/2011 for an Oracle Exadata X2-2 HP Full Rack and a full rack of Netezza TwinFin. The cost to acquire Netezza can be as low as 1/6 of Exadata if a client is acquiring new Oracle database licenses and as low as 1/2 if using existing Oracle database licenses. IBM, the IBM logo, ibm.com, Smarter Planet and the planet icon are trademarks of International Business Machines Corp, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml.

CONTENTS

Windows IT Pro
A PENTON PUBLICATION

APRIL 2011
VOLUME_17 NO_4

COVER STORY

FEATURES

INTERACT

IN EVERY ISSUE

Access articles online at www.windowsitpro.com. Enter the article ID (located at the end of each article) in the InstantDoc ID text box on the home page.

COLUMNS

CONTENTS

PRODUCTS

Windows IT Pro

EDITORIAL

Editorial and Custom Strategy Director
Michele Crockett mcrockett@windowsitpro.com

Editor in Chief
Amy Eisenberg amy@windowsitpro.com

Senior Technical Director
Michael Otey motey@windowsitpro.com

Technical Director
Sean Deuby sdeuby@windowsitpro.com

Senior Technical Analyst
Paul Thurrott news@windowsitpro.com

Industry News Analyst
Jeff James jjames@windowsitpro.com

Custom Group Editorial Director
Dave Bernard dbernard@windowsitpro.com

Developer Content
Anne Grubb agrubb@windowsitpro.com

Exchange & Outlook
Brian Winstead bwinstead@windowsitpro.com

Networking, Storage, Hardware
Jason Bovberg jbovberg@windowsitpro.com

SharePoint
Caroline Marwitz cmarwitz@windowsitpro.com

SQL Server
Megan Keller mkeller@windowsitpro.com

Systems Management, Virtualization, Windows OS
Zac Wiggy zwiggy@windowsitpro.com

Editorial Web Architect
Brian Reinholz breinholz@windowsitpro.com

CONTRIBUTORS

SharePoint and Office Community Editor
Dan Holme danh@intelliem.com

Senior Contributing Editors
David Chernicoff david@windowsitpro.com
Mark Joseph Edwards mje@windowsitpro.com
Kathy Ivens kivens@windowsitpro.com
Mark Minasi mark@minasi.com
Paul Robichaux paul@robichaux.net
Mark Russinovich mark@sysinternals.com

Contributing Editors
Alex K. Angelopoulos aka@mvps.org
Sean Deuby sdeuby@windowsitpro.com
Michael Dragone mike@mikerochip.com
Jeff Fellinge jeff@blackstatic.com
Brett Hill brett@iisanswers.com
Darren Mar-Elia dmarelia@windowsitpro.com
Tony Redmond 12knocksinna@gmail.com
Ed Roth eroth@windowsitpro.com
Eric B. Rux ericbrux@whshelp.com
John Savill john@savilltech.com
William Sheldon bsheldon@interknowlogy.com
Randy Franklin Smith rsmith@montereytechgroup.com
Curt Spanburgh cspanburgh@scg.net
Orin Thomas orin@windowsitpro.com
Douglas Toombs help@toombs.us
Ethan Wilansky ewilansky@windowsitpro.com

ART & PRODUCTION

Production Director
Linda Kirchgesler linda@windowsitpro.com

Senior Graphic Designer
Matt Wiebe matt.wiebe@penton.com

ADVERTISING SALES

Publisher
Peg Miller pmiller@windowsitpro.com

Director, International and Agency Services
Don Knox don.knox@penton.com

Business Development Director
Kerry Gates kerry.gates@penton.com

EMEA Managing Director
Irene Clapham irene.clapham@penton.com

Director of IT Strategy and Partner Alliances
Birdie J. Ghiglione birdie.ghiglione@penton.com
619-442-4064

Online Sales and Marketing Manager
Dina Baird Dina.Baird@penton.com

Key Account Director
Chrissy Ferraro christina.ferraro@penton.com
970-203-2883

Account Executives
Barbara Ritter barbara.ritter@penton.com
858-367-8058
Cass Schulz cassandra.schulz@penton.com
858-357-7649

Client Project Managers
Michelle Andrews 970-613-4964
Kim Eck 970-203-2953

Ad Production Supervisor
Glenda Vaught glenda.vaught@penton.com

MARKETING & CIRCULATION

Customer Service service@windowsitpro.com

IT Group Audience Development Director
Marie Evans marie.evans@penton.com

Marketing Director
Sandy Lang sandy.lang@penton.com

CORPORATE

PENTON

Chief Executive Officer
Sharon Rowlands Sharon.Rowlands@penton.com

Chief Financial Officer/Executive Vice President
Nicola Allais Nicola.Allais@penton.com

TECHNOLOGY GROUP

Senior Vice President, Technology Media Group
Kim Paulsen kpaulsen@windowsitpro.com

WRITING FOR *WINDOWS IT PRO*
Submit queries about topics of importance to Windows managers and systems administrators to articles@windowsitpro.com.

PROGRAM CODE

LIST RENTALS
Contact MeritDirect, 333 Westchester Avenue, White Plains, NY or www.meritdirect.com/penton.

REPRINTS
Diane Madzelonka, Diane.madzelonka@penton.com, 216-931-9268, 888-858-8851

High Availability for MICROSOFT® EXCHANGE Server 2010

IT agility. Your way.

Major re-engineering of Microsoft Exchange Server 2010 offers new functionality for scalability, reliability, and high availability. These changes, however, have consequences for your application infrastructure. Exchange 2010 requires hardware load balancing to ensure high availability because all traffic is now brokered through the Client Access server role.

F5 Application Delivery Networking solutions provide the required core load balancing and additional application-focused features to improve reliability, performance, and security for Exchange 2010 deployments.

Benefits include

- High availability and superior user response.
- Site resilience and disaster recovery.
- Ease of deployment and seamless migration.

Products

BIG-IP® Local Traffic Manager™
BIG-IP® Global Traffic Manager™
BIG-IP® WAN Optimization Module™
BIG-IP® Edge Gateway™
BIG-IP® Message Security Module™

Learn more

Phone: 206.272.5555 or 88888BIGIP
Email: info@f5.com
Web: www.f5.com

f5 IT agility. Your way.

www.f5.com

Windows IT Pro PARTNER

Crockett

"Becoming a SharePoint expert overnight isn't possible, but you can quickly build sufficient skills to help drive your business forward."

Break Through Barriers to SharePoint Success

IT pros hold the key to realizing SharePoint's business optimization potential

One way to bullet-proof your career is to make sure you understand your company's core business processes. IT pros who can support business objectives with cost-effective, efficient IT strategies can directly affect a company's profitability. One of the best examples of IT practices influencing—if not determining—a company's success might be SharePoint deployment and management. No other Microsoft product promises so much to business leaders looking for innovative processes and to IT pros who support those business leaders as well as an army of end users. Businesses use SharePoint most commonly for document management and storage, team collaboration, and web portal and intranet development—all of which are high-profile functions requiring high availability, ease of use, and tight security.

About 44 percent of those who read our *SharePoint Pro Connections* UPDATE email newsletter say that SharePoint is a "very important application" for the organization's operation, and about a third say that it's a mission-critical application. Of those newsletter readers, about half work with SharePoint along with other IT-related responsibilities, and about a quarter focus exclusively on SharePoint. Among our readers overall, about 28 percent currently use SharePoint Server and 38 percent had plans to deploy SharePoint in 2010 or 2011. Given the importance of SharePoint to business processes and the commensurate alignment of IT resources, you'd think that SharePoint would be the magic business bullet.

But as every IT pro knows, the magic is in the implementation. A carefully planned and well-executed SharePoint deployment can yield significant business benefits, including easier information access, better communication, IT time savings, increased end-user productivity, document version control, and savings in web development resources. But IT pros must overcome numerous barriers—including hardware and software deficiencies and organizational problems—to realize those benefits. IT pros grapple with SharePoint implementation problems that include slow system performance, inadequate data protection, inefficient IT management practices required to support the SharePoint environment, problems with migrating data from disparate sources into SharePoint, and slow data backup and recovery processes.

If SharePoint management consumes your day, you've probably successfully navigated through many pitfalls and are now the company IT hero or you've set up your own lucrative consulting practice. But if you just need to solve a pressing SharePoint problem, or you have a short-term SharePoint project to execute, you might find the following resources helpful.

Do you need to get up to speed on SharePoint—fast? If you have a short-term SharePoint project looming and you're starting from scratch from a SharePoint knowledge perspective, check out the SharePoint Collaboration Boot Camp, an intensive workshop taught by Dan Holme. This fast-paced workshop is intended for seasoned IT pros who are new to SharePoint or who are migrating from previous platforms. For a complete workshop description, go to www.devconnections.com/sptour.

Do you need to integrate SharePoint and Outlook for end users? You can extend SharePoint's document management and collaboration functionality by integrating it with Outlook. Check out our *Pocket Guide to Integrating Outlook and SharePoint* at our online store: www.left-brain.com/product.aspx?productid=1283.

Do you need to integrate SharePoint with SQL Server's BI functionality? Stacia Misner can show you how to use SQL Server Reporting Services (SSRS) 2008 R2 in SharePoint integrated mode, which means you'll have only one security model to manage and business users will have only one environment in which to create, find, and share information. Go to www.sqlmag.com and enter InstantDoc ID 129140.

Do you need to migrate to SharePoint 2010? Check out Randy Williams's step-by-step guide at www.sharepointproconnections.com, InstantDoc ID 103467.

Do you need to ensure that your SharePoint implementation is secure? Get a better understanding of the new claims-based security model in SharePoint 2010 at www.sharepointproconnections.com, InstantDoc ID 125108.

Do you need to keep your SharePoint database running smoothly? Matt Ranlett and Brendon Schwartz guide you through standard maintenance tasks to keep SharePoint spinning like a top at www.sharepointproconnections.com, InstantDoc ID 126012.

Do you need quick answers to specific SharePoint problems? Search our SharePoint FAQs at www.sharepointproconnections.com/FAQs.aspx, subscribe to our SharePoint email newsletter at www.windowsitpro.com/email, and subscribe to our bimonthly *SharePoint Pro Connections* magazine—free—at www.sharepointproconnections.com/subscribe.aspx.

Becoming a SharePoint expert overnight isn't possible, but you can quickly build sufficient skills to help drive your business forward. Have you recently had a successful breakthrough with SharePoint? Let me know about your SharePoint experiences at michele.crockett@penton.com.

InstantDoc ID 129782

MICHELE CROCKETT (michele.crockett@penton.com) is editorial strategy director of Penton Media's IT and developer publications, including *DevProConnections, Windows IT Pro, SharePoint Pro Connections, SQL Server Magazine,* and *Connected Planet.*

James

"Perhaps more than any other Microsoft product, SharePoint needs business and IT decision makers to fully understand *why* they're moving to SharePoint before the actual deployment process begins."

Solid Strategy Is Key to SharePoint Success

Deploying SharePoint with proper business goals and objectives in mind is essential to ensuring that your SharePoint rollout reaches its full potential

Everyone has heard the story about the blind men and the elephant: One touches the trunk, one grabs the tail, and another feels the side of the giant pachyderm. All three have a different experience, and all three are technically correct. An elephant is all the things they individually experience—and more.

Such is the case with SharePoint, arguably one of the most versatile and multifunctional platforms to ever emerge from Microsoft. SharePoint can be (and is) accurately described as a collaboration tool, a document repository, a content management system, and a vehicle for developing and maintaining internal and externally facing websites. It's also a social media tool, with blogs, forums, and other community-friendly features.

Despite the success and ubiquity of SharePoint, it can be a major drain on IT resources if it isn't deployed with a clear strategy and business-friendly goals in mind. Perhaps more than any other Microsoft product, SharePoint needs business and IT decision makers to fully understand *why* they're moving to SharePoint before the actual deployment process begins.

I've personally been involved in projects that used SharePoint as an externally facing website publishing system, and the shortcomings and drawbacks of that deployment didn't make it an optimal solution for what it was being used for. After talking to dozens of administrators and users of SharePoint over the past few years, I've come to the conclusion that SharePoint deployments that succeed—and by succeed, I mean achieve the larger goals of the organization, achieve a recognizable return on investment, and are widely embraced by users—seem to have several things in common. We've all seen SharePoint installations that never live up to their full potential, but I'd argue that the implementations that do tend to have three common characteristics. They're deployed with a well-articulated, actionable strategy that meets business goals; the stakeholders are focused on the *why* of deploying SharePoint and not just the *how;* and use and adoption of the SharePoint rollout was supported by training and buy-in from upper management.

1. Have a clear strategy. SharePoint guru Joel Oleson once told me that deploying SharePoint was like reading through one of those *Choose Your Own Adventure* books: You have hundreds of decisions to make during deployment, and some can't be undone by the equivalent of flipping back to an earlier page. It's important to get with all the important stakeholders in your organization—IT, HR, legal, sales, marketing, etc.—and mutually agree on what your SharePoint deployment *will* and *will not* be used for. Doing all your homework and getting agreement between stakeholders might take some time, but that advice will reap significant long-term rewards. I've personally seen and heard about SharePoint deployments that go nowhere, mainly because no one in the organization articulated a clear, well-defined strategy of why SharePoint was needed and why it was being adopted.

2. Focus on the *why*, not the *how*. Once you have a SharePoint deployment and usage strategy in place, your role as a business decision maker or senior IT leader is to help your IT staff communicate to your organization's users why SharePoint is being adopted. If users don't get this information, they might individually make up their own reasons why they should use SharePoint, which could be at odds with the larger goals of your organization.

3. Training, training, and executive investment. We've all seen SharePoint installations that meet the aforementioned criteria but fail where it matters the most: at the user level. SharePoint is an amazing tool, but it's next to worthless if your users don't have the training they need to use the platform properly. Some aspects of SharePoint aren't very intuitive, and a user who could dive into using the latest version of Windows or Office without skipping a beat could find himself drowning in SharePoint. Document check-in and check-out is a vital part of using SharePoint as a document management and collaboration tool, but far too many users have never been instructed how to properly use it. Senior IT executives and other business stakeholders should lead by example by being the first to embrace SharePoint and use it for its intended purpose. If the CEO is editing group business documents using proper check-in and check-out procedures, I guarantee that the rest of the organization will be motivated to learn that function as well.

Have you already embraced SharePoint in a big way in your own organization? Do you have any big-picture SharePoint tips or strategies to share? Send your advice and suggestions to me via email at jeff.james@penton.com, and follow me on Twitter @jeffjames3.

InstantDoc ID 129783

JEFF JAMES (jeff.james@penton.com) is industry news analyst for *Windows IT Pro*. He was previously editor in chief of Microsoft *TechNet* magazine, was an editorial director at the LEGO Company, and has more than 15 years of experience as a technology writer and journalist.

- Career-Killing Cloud?
- Cloud Frauds
- No Certificates

LETTERS@WINDOWSITPRO.COM

Is the Cloud a Career Killer?

I read Michele Crockett's "Securing Your Position in the Cloud" (February 2011, InstantDoc ID 129267). I make my living administering a small company, and I see cloud computing as a potential career killer. So, I was excited to see an article that would address my concerns. After reading it, I feel no better.

Basically, Michele says to research how it works, ask a service provider a few questions, and hold hands for a while until all the services are transferred to the cloud. But here's my story: Last month, a large client of mine approved a huge IT budget to move to an offsite Hyper-V solution. I acted as a trusted source, did some research, answered a few questions, and now—bang!—I'm out of the loop. The guys who took over even have a local tech taking care of all the "cloud appliances" required to connect to the new service. Other than menial jobs around the shop, there will be nothing much for me to do.

What I see working in my community are IT shops buying expensive iSCSI and software-as-a-service NAS boxes, getting into a data center somehow, and moving the service to their hardware—just paying the bills, by the looks of it.

I think cloud computing works wonders for big firms with terminal servers all over the place and locations everywhere. But for the little guy, I'm not so sure. I also think once the price of high-speed NAS comes down, there will be a place for me to move my existing clients' physical servers to virtual servers in their own shops. And that will happen as virtual computing becomes more robust than physical servers.

—Craig Musgrove

Publishing Updates Without Certificates

Russell Smith's article, "Publishing Third-Party Updates to WSUS" (February 2011, InstantDoc ID 129241) discusses using the open-source Local Update Publisher to publish third-party updates to WSUS—a capability I've been wanting for a long time. Thank you for this article! Maybe this capability has been around for a while, but I didn't know about it. Either way, *Windows IT Pro* has given me nuggets of gold like this quite often.

In an environment without a PKI server, can this be done without certificates? Creating the self-signed certificate as you describe is no problem. Getting it onto the clients is. Any advice? Can this be done without certificates? (I'm trying that now.)

—Dan Wakeman

I'm glad that you find the material in Windows IT Pro *useful. To answer your question, there's no supported way to do this without certificates. But self-signed certificates, or certificates purchased from a third-party CA, can be distributed to clients using Group Policy or the Certutil/Certmgr command-line tools.*

—*Russell Smith*

Cloud Frauds

Although I agree with Jeff James' assessment of cloud computing ("Why IT Is Moving to the Cloud," February 2011, InstantDoc ID 129285), I think there should be some cautions regarding many companies who have simply rebranded themselves to join the revolution. These companies do a lot to hurt the concept of cloud computing while trying to ride the "wave" for free. Those who are serious about cloud computing waste a lot of time and money verifying who they're really dealing with and what their actual capabilities are.

—Eugene O'Neal

InstantDoc ID 129760

Windows IT Pro welcomes feedback about the magazine. Send comments to letters@windowsitpro.com, and include your full name, email address, and daytime phone number. We edit all letters and replies for style, length, and clarity.

Thurrott

"SP1 provides new features to Windows Server 2008 R2 and is as big a leap over R2 as was R2 over Windows Server 2008, at least from a Hyper-V perspective."

Diving into SP1 for Windows 7 and Windows Server 2008 R2 and More Windows Phone 7 Updates

Although I typically address a wide range of products, technologies, and trends in this column, this month two big updates warrant our undivided attention. We've touched on the first service pack for Windows 7 and Windows Server 2008 R2 in previous issues. But I'd like to take a deeper look because SP1 is now available publicly and is ready to be deployed in production environments. Then I'll examine how Microsoft intends to update Windows Phone 7 in 2011.

SP1 for Windows 7 and Server 2008 R2

Microsoft released SP1 to manufacturing on February 10, 2011 and delivered the code to its hardware partners the same day. Releases to volume license customers and TechNet and MSDN subscribers followed on February 16, and the service pack was publicly released via Microsoft Update, manual download, and other avenues on February 22. That latter date was also the day on which server makers could begin selling new machines with the SP1 code preinstalled.

SP1 services both Windows 7 and Windows Server 2008 R2. (This follows the development of Windows Vista and Windows Server 2008, which were also derived from the same code base as each other and are serviced by their own service packs.) But though each product utilizes the same service pack, there are two different versions of SP1: one that targets 32-bit versions of the OS and one that targets 64-bit (x64) versions.

For Windows 7, SP1 is a minor update consisting mostly of previously released fixes and a handful of minor functional updates. However, Windows Server 2008 R2 SP1 is considerably more dramatic. It, too, supplies a number of mostly previously released fixes. But it also includes two major new features, Dynamic Memory and RemoteFX, both of which significantly enhance the capabilities of the Hyper-V virtualization platform in Windows Server.

Impact on Windows 7

Because business customers have traditionally waited for the first service pack update to any major Microsoft OS before deploying, the software giant communicated that waiting wouldn't be necessary with Windows 7. And sure enough, business customers have demonstrated a desire to migrate to Windows 7 much earlier in its life cycle than has been the case with previous Windows versions. This can't be attributed to the quality of Windows 7, however. Instead, it's more likely that the several-year gap between Windows XP and Windows 7—marked by the release of the disastrous Vista—is the real cause.

But whatever the reason, SP1 doesn't change Windows 7 in any major ways, though its minor changes could be important if they address issues you've had. These changes include

Remote Desktop Services update. This is required for the new Server 2008 R2 feature, RemoteFX, described below.

Better support for third-party federation services. With this update, Windows 7 now supports services that utilize the WS-Federation passive profile protocol.

Improved HDMI audio device performance. SP1 fixes a bug in Windows 7 where a small percentage of users experienced a disruption of audio over HDMI after a reboot.

Minor XPS document fixes. For the rare case where an XPS document contains both portrait and landscape pages, SP1 fixes a bug that prevented correctly-formatted printing.

Hotfixes and other bug fixes. Like all service packs, SP1 also contains an aggregation of previously released and new hotfixes and other bug fixes. According to Microsoft documentation, SP1 contains well over 600 individual hotfixes for both Windows 7 and Server 2008 R2.

Impact on Windows Server 2008 R2

SP1 provides two major new features to Server 2008 R2 and is, in some ways, as big a leap over R2 as was R2 over the original shipping version of Server 2008, at least from a Hyper-V perspective. These updates bring the Hyper-V virtualization platform in Server 2008 R2 closer to the capabilities of the VMware competition while adding a unique virtual desktop infrastructure (VDI) capability that will appeal to larger businesses with tightly-controlled client environments.

Microsoft introduced Hyper-V with Server 2008 in 2008, providing Windows Server with a type-1 hypervisor-based virtualization solution. (In fact, the initial version of the OS shipped with a prerelease version of Hyper-V.) In Server 2008 R2, Hyper-V was augmented with live migration capabilities, improved performance, and core parking functionality. In Server 2008 R2 SP1, Hyper-V

is improved yet again with two major new features:

Dynamic Memory. This feature provides the server with a way to automatically and efficiently manage the memory allotted to running virtual machines (VMs), giving it the ability to dole out memory based on demand. This expands the number of VMs a given hardware server can support simultaneously, or what Microsoft calls increasing VM density. And it does so without any performance sacrifices.

Dynamic Memory works with 32-bit or 64-bit Vista, Windows 7, Windows Server 2003, Server 2008, and Server 2008 R2 (64-bit only) clients (or what Hyper-V calls child partitions). More specifically, it can't work with XP because of limitations in that version of the OS. (Support for XP is also being phased out, finally.)

With regards to server density, on the initial shipping version of Server 2008, a server with 96GB of RAM could handle about 85 Windows 7 VMs, or 85 to 120 XP VMs (where each client is configured with 1GB of RAM; these figures are from Microsoft). Back then, the big scaling constraints were physical in nature: The servers of the day just couldn't be expanded with more RAM. Today, with Server 2008 R2 SP1, just enabling Dynamic Memory increases VM density by an average of 40 percent (again, according to Microsoft). Microsoft was able to effectively run 120 Windows 7 (32-bit) clients on a server with just 87GB of RAM, for example, and that was without maxing out resources anywhere.

One interesting performance note: For client versions of Windows, 32-bit versions of Windows 7 offer dramatically better server density than 64-bit versions due to their lower runtime memory requirements. According to Microsoft, typical business workloads require about 540MB of RAM per 32-bit Windows 7 VM compared to 720MB for 64-bit versions.

RemoteFX. This feature adds server-side hardware-based graphics acceleration for VMs based on Windows 7 Enterprise and Windows 7 Ultimate. It virtualizes the server's GPU and makes it available to supported clients. As such, it's most useful in VDI environments, where thin clients are used on desktops and most of the processing investments are in the data center.

Obviously, servers with GPUs aren't particularly common today, but that is changing with the advent of SP1. And for those customers who choose the VDI route over, say, optimized Windows 7 desktops, RemoteFX provides a way to gain back part of the client-rendering punch normally sacrificed by such a move. This makes it possible to utilize DirectX and Direct3D applications, including media players and the like. (OpenGL isn't explicitly supported, but some OpenGL apps that work through DirectX will work.)

RemoteFX works as you'd expect: The VM believes it's talking directly to a GPU and not indirectly communicating to a GPU located on the server. This is full hardware-accelerated rendering, so applications like Internet Explorer 9 that can utilize the GPU on a Windows 7 PC can also do so in a VM now too.

As part of the RemoteFX feature set, Remote Desktop Services (RDS) has been updated to 7.1 thanks to some performance improvements around encoding, decoding, and session management. And Hyper-V can now do USB redirection, providing VMs with the ability to interact with connected devices such as cameras and microphones.

Acquiring SP1

As with any service pack, SP1 is available via a variety of Microsoft delivery methods. The standalone package, which includes support for five languages and will be of most interest to system administrators, ranges in size from about 300MB for the 32-bit version to about 535MB for the 64-bit version.

Those who interactively install SP1 via Windows Update or Microsoft's other automated updating services will experience different download sizes and installation times based on how up to date their OS is at the time of install. This is because the SP1 interactive installer examines the contents of the system and ensures that only new or changed components are downloaded and installed. According to Microsoft, the download size could range from a low of about 23MB (32-bit Windows 7) to as much as 50MB (64-bit Windows Server 2008 R2).

Microsoft will also be supplying integrated DVDs and ISOs that combine various versions of Windows 7 and Server 2008 R2 with the SP1 code. These will be distributed through the normal channels, including retail, volume licensing, MSDN, and TechNet.

More About 2011 Software Updates for Windows Phone 7

At the Mobile World Congress trade show in February, Microsoft revealed its plans for updating Windows Phone 7 in 2011, specifying some of the contents of two eagerly awaited upgrades to its fledgling mobile OS. The first, code-named No Donuts ("NoDo"), will ship by the middle of March and provide such features as copy and paste, application launch performance improvements, and better Marketplace search. This update will also enable CDMA support, triggering the launch of Windows Phone 7 devices on the Sprint and Verizon wireless networks.

The second update, expected in the second half of 2011, is a major update that will dramatically enhance the Windows Phone user experience. My sources tell me it should ship at roughly GA + 1 ("general availability plus one," or one year after the launch of the original Windows Phone).

Among the changes coming in this second update, code-named Mango, are Twitter integration into the People hub (similar to today's Facebook integration), an upgrade to IE 9 Mobile (which will be hardware-accelerated just like its PC-based brethren), Windows Live SkyDrive integration for Office document sync to the cloud, and true multitasking for third-party applications.

These updates can't come fast enough in my opinion, and I do expect Microsoft to close the enterprise feature gap with Windows Mobile by the end of 2011 as well. Stay tuned to the SuperSite for Windows (www.winsupersite.com) for more information about the future of Windows Phone.

InstantDoc ID 129741

PAUL THURROTT (thurrott@windowsitpro.com) is the senior technical analyst for *Windows IT Pro*. He writes a weekly editorial for *Windows IT Pro UPDATE* (www.windowsitpro.com/email) and a daily Windows news and information newsletter called *WinInfo Daily UPDATE* (www.wininformant.com).

Get 71% greater performance than with Oracle WebLogic (and pay only for cores you use).

Head-to-head performance is no contest. IBM WebSphere® Application Server on Power Systems™ simply outperforms. Add in fair pricing—IBM doesn't charge you for cores you aren't using—and the case is closed. Saving 57% on first-year licensing and support? We'll mention that in passing so it doesn't look like we're rubbing it in. Find out more about IBM advantages.

ibm.com/facts

PERFORMANCE comparison based on SPECjEnterprise2010 results from www.spec.org as of 2/10/2011 and compares performance per core of the WebSphere Application Server V7 on IBM Power 730 Express and DB2 9.7 on IBM BladeCenter PS701 Express result of 4,062.38 EjOPS on 16 cores against Oracle WebLogic Server Standard Edition Release 10.3.3 on Oracle SPARC T3-4 score of 9,456.28 EjOPS on 64 cores. SPEC and SPECjEnterprise are registered trademarks of the Standard Performance Evaluation Corporation. SAVINGS based on publicly available information as of 2/10/2011 comparing Oracle WebLogic Server Enterprise Edition to IBM WebSphere Application Server Network Deployment, both on an IBM Power 730 Express server (2 chips, 8 cores each). IBM, the IBM logo, ibm.com, Power Systems, WebSphere, Smarter Planet and the planet icon are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml.

THE CONVERSATION BEGINS HERE

SharePoint CONNECTIONS Coast to Coast Tour + Microsoft SharePoint BOOTCAMP

COMING TO A CITY NEAR YOU IN 2011!

DIVE INTO SHAREPOINT 2010 WITH MICROSOFT AND SHAREPOINT INDUSTRY EXPERTS.

Register now for the developer and IT pro bootcamps. **SPACE IS LIMITED!**

EARLY BIRD fee: $499

REGULAR fee: $599

REGISTER EARLY

The first 100 developers to register for the ***SharePoint Coast-to-Coast Tour*** in each city get into the hands-on *Microsoft SharePoint 2010 Development Bootcamp* for **FREE**!

A SAMPLING OF SPEAKERS

MICHAEL NOEL
CONVERGENT COMPUTING

STEVE FOX
MICROSOFT

DAN HOLME
INTELLIEM, INC.

RICHARD TAYLOR
IGOTIT TECHNICAL SERVICES

TODD BAGINSKI
FRESH TRACKS CONSULTING, LLC

MATT MCDERMOTT
ABLEBLUE

SCOT HILLIER
SCOT HILLIER TECHNICAL SOLUTIONS, LLC

ASIF REHMANI
SHAREPOINT ELEARNING.COM

CHRIS GIVENS
ARCHITECTING CONNECTED SYSTEMS

PAUL STUBBS
MICROSOFT

DARRIN BISHOP
KNOWLEDGELAKE, INC

ROBERT L. BOGUE
THOR PROJECTS

ANDREW CONNELL
CRITICAL PATH TRAINING, LLC

RANDY WILLIAMS
SYNERGY CORPORATE TECHNOLOGIES

TO REGISTER: DevConnections.com/SPTour 800.438.6720

Minasi

"It's a bit kludgy, but it's a viable replacement for SteadyState."

Replicating SteadyState in Windows 7

Yes, you can create physical snapshots in the latest Windows OS!

One of my favorite virtual machine (VM) features is *snapshots*. Want to test something messy on your VM but don't want to risk the necessity to rebuild the system afterward? Just take a snapshot!

Snapshots aren't just a convenience. Many educational institutions support *physical* (not virtual) student machines that need to be restored to their like-new state after every class. Microsoft used to offer a free tool called SteadyState that did just that thing for Windows Vista and Windows XP, but the tool is missing from Windows 7. So, this month I want to show you how to create your own homemade SteadyState functionality for Windows 7 Enterprise and Ultimate editions.

You can implement physical snapshot/recovery capabilities for a physical Window 7 system by combining two technologies from past columns: Windows 7 Enterprise/Ultimate's boot-from-VHD capability ("Booting Windows 7 Enterprise or Ultimate from a VHD File," InstantDoc ID 129377), and those OSs' ability to natively support a VHD type called a *differencing disk* ("Diskpart Takes Snapshots of Physical and Virtual Systems," InstantDoc ID 125233).

First, get your desktop on a VHD. You can use the procedure explained in "Creating a Bootable VHD" (InstantDoc ID 129194) to set up a Windows 7 system as a bootable VHD. Call that file baseimage.vhd. (Be sure to make baseimage.vhd an *expandable* VHD—not fixed-size—or you won't be able to create a snapshot.) You'll want to install baseimage.vhd as a second boot option on your computer, as you read in "Creating a Bootable VHD," so copy that file to your computer and put it in a folder named C:\VHDs.

Second, using the Bcdedit commands you saw in "Booting Windows 7 Enterprise or Ultimate from a VHD File," configure your system to boot from baseimage.vhd. At that point, use the new option to boot from baseimage.vhd and do whatever you need to get your system just the way you'd like it, prior to *snapshotting* it.

Now, it's time to create the snapshot. You want to be able to boot from baseimage.vhd while at the same time being able to undo any changes to baseimage.vhd, so you'll tell Windows 7 to leave baseimage.vhd unchanged and instead write any changes to a *different* VHD by creating a new VHD that is neither fixed nor expandable but is instead a *differencing* VHD associated with baseimage.vhd. Windows refers to the original VHD as the *parent* VHD and the new differencing VHD as the *child* VHD. Windows won't let you create a child VHD to an in-use parent VHD, however, so before you can create a child for baseimage.vhd, you need to reboot to the original Windows image on the computer's C drive to free up baseimage.vhd. Now, open Diskpart and type

```
create vdisk file=c:\vhds\snapshot1.vhd parent=c:\vhds\
  baseimage.vhd
```

One more step, and you're in business: Using the Bcdedit commands that you used to enable your system to boot from baseimage.vhd (again, your system's second boot option), create a *third* boot option, booting this time from C:\vhds\snapshot1.vhd.

Boot from that third option, and from this point on baseimage.vhd remains unchanged, and all changes go into snapshot1.vhd. You can then modify your system however you like and, when you want to return things to their pre-snapshotted state, just reboot into the first of your three boot options, open an elevated command prompt, and type

```
del c:\vhds\snapshot1.vhd
diskpart
create vdisk file=c:\vhds\snapshot1.vhd parent=c:\vhds\
  baseimage.vhd
exit
```

After a reboot into the third option, your system is back to its base state. As before, any new changes will go into snapshot1.vhd, and if you ever want to return to the "pristine state," just recreate the above commands to continue to use the third boot option.

If you want to update baseimage.vhd in some way, you have two choices. Either boot to the second option and do your maintenance (the second option writes all changes to baseimage.vhd) or, if you like the changes that you've wrought to snapshot1.vhd and want to keep them, boot from the first option and merge snapshot1.vhd into baseimage.vhd like so:

```
diskpart
select vdisk file=c:\vhds\snapshot1.vhd
merge vdisk depth=1
exit
del c:\vhds\snapshot1.vhd
```

Yes, I know, it's a bit kludgy. But try it out. I think you'll see that it's a viable replacement for SteadyState. As to the kludgy part? I'll offer a few ideas to smooth that out next month.

InstantDoc ID 129192

MARK MINASI (www.minasi.com/gethelp) is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 25 books.

Top 10 Free Tools for System Administrators

Audit Active Directory and file servers, detect inactive users, block USB devices, and more – for free

The following freeware tools by Windows IT Pro Community Choice Awards finalist NetWrix Corporation can save you a lot of time and make your network more efficient – at absolutely no cost. All of these tools also have advanced commercial editions with additional features, but the freeware editions will not expire and will not stop working when you urgently need them.

1. Active Directory Change Reporter (Windows IT Pro Sep'09: InstantDoc ID 102446, TechRepublic: www.tinyurl.com/4az79au)—This simple auditing tool keeps tabs on what's going on inside your Active Directory. The Windows IT Pro 2010 Community Choice and Editors' Best Award-winner tracks changes to users, groups, OUs, and all other types of AD objects, sending detailed daily reports with lists of changes. Download link: www.tinyurl.com/47l82dy

2. Privileged Account Manager (SC Magazine: www.tinyurl.com/4pu9gqa) —This product maintains a repository of privileged user accounts (such as Administrator, root, service accounts etc) in Active Directory, servers, and other systems, providing a secure web-based portal for role-based access and automatic maintenance of shared administrative user accounts. The Privileged Account Manager can automatically generate strong passwords at specified intervals (e.g. every 30 days) and synchronize password changes on all target systems (for example, change service account password in Active Directory and update service credentials). Download link: www.tinyurl.com/4jtavfc

3. USB Blocker (Windows IT Pro Nov'09: InstantDoc ID 102860)—The increasing mobility of flash drives, MP3 players, cell phones and iPods makes the threat of data theft greater than ever, and with a couple clicks of the mouse, this aptly-named tool blocks unauthorized usage of removable media via USB ports. USB Blocker hardens end point security by preventing the spread of harmful malware and restricting the transfer of confidential information. Download link: www.tinyurl.com/4dtn5qk

4. Password Expiration Notifier (Redmond Magazine Feb'09, 4sysops: www.tinyurl.com/4ost9py)—This tool automatically reminds users to change their passwords before they expire, helping keep helpdesk administrators safe from password reset calls. It works nicely for users who don't log on interactively and, thus, never receive standard password change reminders at log on time (VPN and OWA). Download link: www.tinyurl.com/49aox36

5. Inactive Users Tracker (TechRepublic: www.tinyurl.com/4zyl552)—This tool tracks down inactive user accounts (e.g., terminated employees) so you can easily disable them, or even remove them entirely, thus eliminating potential security holes. The tool sends reports on a regular schedule, showing what accounts have been inactive for a configurable period of time (e.g., 2 months). Download link: www.tinyurl.com/4cwflbp

6. File Server Change Reporter (4sysops.com: www.tinyurl.com/4pe435y)—This is a must-have tool for auditing file servers and appliances. The tool detects changes made to files, folders and permissions, and tracks newly created and deleted files. The tool is useful for detecting mistakenly deleted files and it allows quick backup recovery of accidental changes. Download link: www.tinyurl.com/47qgag9

7. Active Directory Object Restore Wizard (Windows IT Pro: www.tinyurl.com/4eb454e)—This tool can save the day if someone accidentally (or intentionally) deleted important Active Directory objects. It provides granular object-level, and even attribute-level restore capabilities that allow quick rollbacks of unwanted changes (e.g., mistakenly deleted users, modified group memberships, etc). Download link: www.tinyurl.com/4nfy9r9

8. VMware Change Reporter (TechTarget/SearchVirtualDesktop: www.tinyurl.com/4odfghv) — If you don't know what is being changed by your colleagues in the VMware infrastructure, it's very easy to get lost and miss changes that can affect the things for which you are responsible. This tool tracks and reports configuration changes in VMware Virtual Center settings and permissions. Download link: www.tinyurl.com/4m6ybyb

9. Windows Service Monitor (WindowsReference.com: www.tinyurl.com/4esgmoc)—This very simple monitoring tool alerts you when some Windows service accidentally stops on one of your servers. The 2010 Windows IT Pro Community Choice and Editor's Best Award-winning tool also detects services that fail to start at boot time, which can happen, for example, with Microsoft Exchange. Download link: www.tinyurl.com/4ud8oct

10. Disk Space Monitor (MS TechNet Magazine Sep'09: www.tinyurl.com/4tfpz2r)— Even with today's terabyte-large hard drives, server disk space tends to run out quickly and unexpectedly. This simple monitoring tool will send you daily reports regarding all servers that are running low on disk space, below the configurable threshold. Download link: www.tinyurl.com/4ut29sz

Otey

"The Windows ecosystem is so mature that free tools abound, and they address the majority of the problems you're likely to run into."

Free Tools for Managing Windows

Make network troubleshooting simple with this assortment of freeware

Free tools are definitely one on my favorite things, and in fact they're one of the best things about the Windows ecosystem: It's so mature and ubiquitous that free tools abound, and they address the majority of the problems and troubleshooting situations you're likely to run into. There are so many free tools that you can't possibly list them all, but here are ten tools that I've recently found to be useful.

10 **PC Inspector File Recovery**—Recently I needed to undelete a folder of photos that was accidently deleted, and the Recycle Bin had been emptied. After trying and discarding a number of free tools, I ran across PC File Inspector from CONVAR at www.pcinspector.de/Default.htm?language=1. This utility isn't supported for Windows Vista, but it ran fine for me in Windows 7.

9 **SDelete**—If you want to delete highly secure or sensitive files, a great tool is Sysinternals SDelete. Instead of merely deleting the file's directory entries, SDelete writes over all of the file's existing data, making it impossible to recover. You can download SDelete from the Windows Sysinternals website at technet.microsoft.com/en-us/sysinternals/bb897443.

8 **FreeFileSync**—Another common task is synchronizing files and folders, which is handy for moving a set of files to your laptop for travel or for comparing different sets of files and folders on your local system. FreeFileSync can compare files and folders and optimally synchronize them. You can download FreeFileSync from SourceForge at sourceforge.net/projects/freefilesync.

7 **ExamDiff**—If you want to compare the contents of files, one capable freeware tool that lets you do so is ExamDiff. ExamDiff provides a graphical interface that immediately shows the differences between two files. You can find ExamDiff on PrestoSoft's website at www.prestosoft.com/edp_examdiff.asp.

6 **7-Zip**—Windows's built-in ability to deal with .zip files is limited. The free open-source 7-Zip tool lets you work with many more file compression formats, including the Linux TAR format. There are 32-bit and 64-bit versions of 7-Zip, and it's completely integrated into Windows Explorer. You can download this tool from the 7-Zip page at www.7-zip.org.

5 **Remote Desktop Manager**—If you're like me, you use a lot of remote desktop sessions to manage the different servers on your network. This is doubly true when many of your systems are running in virtual machines (VMs). Remote Desktop Manager helps you keep all your RDP sessions together. It also supports both RDP and Virtual Network Computing (VNC) remote connections. It's available from Devolutions at remotedesktopmanager.com/remotedesktopmanager/Home.aspx.

4 **WinDirStat**—Sometimes it's difficult to tell which files and folders are consuming the most disk space on your system. WinDirStat is a SourceForge tool that graphically displays disk utilization. It shows a typical navigational tree and also shows the size of each directory as well as a graphical representation of the space consumed by each file. You can download this tool from the WinDirStat page at windirstat.info.

3 **AutoRuns**—Another Sysinternals tool that I often use is AutoRuns, which is a valuable tuning and troubleshooting tool. Like the built-in Windows msconfig tool, AutoRuns shows you what's running when your system starts up. However, AutoRuns takes this analysis to a whole different level, showing all the registry entries, tasks, services, and other boot locations that can be used to automatically run programs at boot up. You can download AutoRuns from Sysinternals at technet.microsoft.com/sysinternals/bb963902.

2 **DNSDataView**—If you're like me and you can never remember all the commands to run Nslookup, you might want to check out DNSDataView. Like Nslookup, DNSDataView lets you retrieve and view DNS records for a specified domain. However, unlike Nslookup, DNSDataView provides a graphical interface. You can get DNSDataView from NirSoft at www.nirsoft.net/utils/dns_records_viewer.html.

1 **Lansweeper**—One of my favorite utilities for generating an inventory of my networked system is Lansweeper. There are server versions of Lansweeper, but the freeware version is capable of producing a basic inventory of all the systems on your network. You can run reports showing different types of servers and clients on your network. The freeware version is limited to scanning one domain. You can download this tool from the Lansweeper page at www.lansweeper.com.

InstantDoc ID 129626

MICHAEL OTEY (motey@windowsitpro.com) is senior technical director for *Windows IT Pro* and *SQL Server Magazine* and author of *Microsoft SQL Server 2008 New Features* (Osborne/McGraw-Hill).

Deuby

"We've developed our IT authorization and authentication systems based on methods we use in the real world."

Border Crossings in the Identity Realm

Our digital authorization and authentication systems are based on methods we've mastered in the physical world

Nothing makes an identity professional think about his work more than crossing into another country. I had the opportunity to consider this truism six times over the holiday break, taking advantage of geography by using Canada as a shortcut between different family locations.

I grew up in Michigan, and my wife grew up in New York. The quickest route between these two locations is through the southern tip of Ontario. Between flying into Toronto to avoid the holiday mess at Chicago O'Hare and driving between the Michigan and New York destinations, I became fairly familiar with each country's border procedures. While I sat in line at the US border, German Shepherds sniffing around the car, I thought it would be interesting to compare a couple of examples of how authentication works in the physical world with their digital counterparts, and how an emerging class of applications mimics and perhaps improves on what's being done in the physical world.

As in the physical world, authentication in the digital realm can involve only a simple password or it can use complex multiple factors such as one-time passwords, biometric scanners, time limitations, and location restrictions.

The Physical Realm

First, let's pick apart what's happening in physical authentication from an identity professional's point of view. In day-to-day authentication, a retail clerk asks to see your driver's license—for example, to see if you're old enough to buy alcohol. (This hasn't happened to me in *way* too long, by the way.) The clerk looks at the license with varying degrees of scrutiny to see if the license appears to be genuine, the photo matches you, the description matches you, and the signature on the license matches your signature in front of them. Most people, however, simply look at the photo and confirm a pattern match with the person standing in front of them.

What is a driver's license, anyway? It's a token. This driver's license "token" has attributes such as a photo, height, weight, and date of birth. It has an expiration date. It's issued by an authority—the state—that certifies the validity of the values of these attributes. That certifying authority requires a variety of supporting documents, as the US driver's license is accepted as a means of establishing identity. Its scope as identity credentials, however, is limited to the United States because proof of US citizenship isn't required to get a driver's license. This is very similar in structure to a Kerberos ticket used in Active Directory (AD) authentication and authorization, or a Security Assertion Markup Language (SAML) token used in claims-based authentication for Internet single sign-on (SSO). Both contain a set of attributes with values, and both are issued by a certifying authority.

A passport is also a token, with similar attributes. The primary difference is that a passport's scope is international because it establishes nationality as well as basic identity characteristics. The certifying authority is the US government, and the document requirements to be issued a passport are more stringent than those of a driver's license. Unlike a driver's license, it also has the ability to carry updates by other certifying authorities (the visa section where passport control puts its stamp) after the passport has been issued.

What happens when you drive up to a border crossing into the United States? The obvious checks are confirming that your passport is valid and matches your description, and checking the car's license for ownership and any outstanding warrants. I'm not a homeland security expert, but it's safe to say that these checks are a small part of the checks that are done. For example, I recently learned of a fellow who was pulled aside by Customs coming into the United States because border security had detected the residue of a radiological agent! (The man had undergone a physical that involved radiology.)

The most important action the border agent performs, however, is to ask you questions and watch your behavior as you answer those questions. After all, not too many people drive up to the border with something as obvious as a stolen car; behavioral questioning can help expose inconsistencies and falsehoods that simple passport authentication doesn't expose. As in the physical world, authentication in the digital realm can involve only a simple password or it can use complex multiple factors such as one-time passwords, biometric scanners, time limitations, and location restrictions.

A border crossing involves both authentication and authorization. Authentication determines whether you're really who you say you are. Authorization determines what resources you're allowed to access and at what level. Once your identity is verified at the border, there's a chance you could find yourself on a watch list that denies your entrance into the country; authorization in this physical case is pretty much binary; you're either allowed in or you aren't. If you're authorized, there's no restriction to shop only at certain stores in certain states. Apart from the obvious constitutional and legal reasons, this is because national passport authentication systems aren't integrated with commercial systems.

The Digital Realm

In the physical world, many tasks come with some basic contextual authorization. Two examples are age checks for buying alcohol and for purchasing R- or NC-17-rated movies. Most of the time, a clerk authenticates you by simply looking at your driver's license to be reasonably sure you're who you say you are. In the digital world, the photo check doesn't happen—one reason that fraud is easier to perform there.

Unlike the casual driver's license check, almost all credit card transactions today are checked in real time for authorization against the credit card's issuer. Thanks to mobile wireless point-of-sale devices, I've had my credit card checked everywhere from art shows in open fields to small businesses in rural Bali.

When every transaction depends on authorization, however, the entire authorization infrastructure must be robust and fault-tolerant because there are real consequences for the user trying to access resources at the far end of the process. For example, I've had my MasterCard account closed while in remote locations due to a "merchant security compromise"—someone hacked a retail company where I'd shopped, and the retailer kept my credit card information without my permission.

This is a real problem if you travel internationally. These credit card authorization systems are somewhat context-sensitive, too, thanks to anti-fraud technology; if most of your credit card purchases are in the United States, and you purchase something in England, don't be surprised if you get a robocall from your credit card issuer requesting that you verify the purchase.

SIEM Systems

These systems are a simple example of a category of software that takes this practice to a sophisticated level. It's called security information event monitoring (SIEM). SIEM software has two major functions, and the SIEM acronym contains both. The first function, security information management (SIM), is to collect information from event logs across potentially hundreds or thousands of systems into a database and provide intelligence and reporting about these systems. This is usually done for compliance or regulatory reasons.

The second function is, to me, the more interesting use. Security event management (SEM) goes beyond simply collecting and reporting on log information; this type of software actively monitors event logs in real time, intelligently analyzes their output, and takes action based on the rules the SEM administrator establishes—for example, to flag or lock out the perceived threat to the system. Have you ever logged on to Facebook or Gmail from a different location than you usually do? You're prompted for a security question. This is because SIEM software in these services has detected that your session is located at a different IP address than your previous ones.

Thanks to mobile wireless point-of-sale devices, I've had my credit card checked everywhere from art shows in open fields to small businesses in rural Bali.

In my experience, if I log on from a significantly different geographic location—say, logging on in Canada when I'm usually in Texas—the next time I log on at my usual IP address range, Gmail displays a warning that there's been an unusual logon for my account and asks me to approve it. Financial companies are in the lead for adopting SIEM software, first for SIM purposes but now for SEM and fraud detection.

A SIEM system has the ability to collect and correlate potentially millions of records from a wide range of network devices—from network switches to endpoint protection devices to line-of-business (LOB) systems. They can detect insider threats such as large amounts of printing after hours, large amounts of email or large attachments to personal email addresses, or unusual system audit log clearing. They can also detect malicious intrusions from external hackers. If so configured, the SIEM package can lock out the user's account. Though these systems aren't cheap to purchase or to implement, they can pay for themselves in minutes if they detect, warn, and perhaps thwart malicious insider action, fraud, or external attacks.

If you're interested in learning more about SIEM products, Gartner published a "Magic Quadrant for Security Information and Event Management" report (www.gartner.com/technology/media-products/reprints/nitrosecurity/article1/article1.html) that describes the various SIEM vendors and their relative strengths and weaknesses. SIEM software is similar to a border agent's behavioral questioning; by correlating various bits of information gathered from many sources, the agent uses his or her training to flag individuals for further investigation.

Lessons Learned

We've developed our IT authorization and authentication systems based on methods we use in the real world. The online world can perform faster and more efficiently than humans, but only recently has it begun to develop the security intelligence that the human element has long provided us. At the same time, human security systems have come to depend more and more on IT security systems to augment their methods. And we need all these methods to combat increasingly sophisticated attempts to invade our systems.

InstantDoc ID 129583

SEAN DEUBY (sean@windowsitpro.com) is technical director for *Windows IT Pro* and *SQL Server Magazine*, and former technical lead of Intel's core directory services team. He's been a directory services MVP since 2004.

READER TO READER

Find PCs Running Out of Disk Space

When I was working in the manufacturing realm, we had some older computers that were still alive and kicking. For the most part, they performed their tasks perfectly, except when the 4GB hard disk ran out of space. Unlike our servers, our PCs didn't have a handy utility that reported incidents like that. However, I still needed to know when a PC was running out of space before the system went down, so I created my own solution, which runs as little code as possible on the PC.

Joe Vogelsong

Here's what I did. First, I added the code in Check_PC_Free_Space.vbs (see Listing 1) to our logon script. When a computer has less than 500MB of free disk space, this code writes the computer's name, the number of bytes free, and the date to a log file. The log file is in a shared folder on a file server. I typically create a shared folder for tasks like this, giving the folder separate permissions. In this case, I granted the Domain Users and Domain Computers groups write access to the shared folder. If you have nosy users, you can hide the shared folder by adding a dollar sign ($) to the end of its name.

Next, I scheduled a task on the server where the log file resides. The scheduled task runs the Log_File_Update_Check.vbs script every four hours. As Listing 2 shows, this script uses VBScript's DateDiff function to determine whether the log file has been updated. If it has been updated, the script executes a batch file that emails the log file to me.

Listing 3 shows the batch file, Email_Log_File.bat, which uses Blat to send the email. The email notification is the most important part of the solution. Without this notification, it would be easy to forget to review the log file regularly. The email lets you know about the problem within a few hours of it being discovered. Thus, the problem can usually be fixed before a Help desk call is made.

If you want to try this solution, you can download the code in the listings by going to www.windowsitpro.com, entering 129546 in the InstantDoc ID box, clicking Go, then clicking the *Download the Code Here* button. You'll also need to download Blat (www.blat.net) if you don't already have it. Blat is great for emailing notifications and various types of files and is pretty simple to use.

—Joe Vogelsong, network administrator

InstantDoc ID 129546

Listing 1: Check_PC_Free_Space.vbs

```
strComputer = "."
Const ForAppending = 8
Dim myDate
myDate = Date()
Set WshNetwork = WScript.CreateObject("WScript.Network")
Set objFSO = CreateObject("Scripting.FileSystemObject")
Set objTextFile = objFSO.OpenTextFile _
  ("\\NetworkShare\LowDriveSpace.log", ForAppending, True)
Set objWMIService = GetObject("winmgmts:" _
  & "{impersonationLevel=impersonate}!\\" & strComputer & "\root\cimv2")
Set colDisks = objWMIService.ExecQuery _
  ("Select * from Win32_LogicalDisk Where DeviceID = 'C:'")
For Each objDisk in colDisks
  If objDisk.FreeSpace < 524288000 Then
    objTextFile.WriteLine (WshNetwork.ComputerName & " , " _
    & objDisk.FreeSpace & " , " & myDate )
  End If
Next

objTextFile.Close
WScript.Quit
```

Listing 2: Log_File_Update_Check.vbs

```
strFile = "G:\Scripting\Reports\
  LowDriveSpace.log"
Set objFSO = CreateObject("Scripting
  .FileSystemObject")
set objFile = objFSO.GetFile(strFile)
Set objShell = WScript
  .CreateObject("WScript.Shell")
dFileModDate = objFSO.GetFile(strFile)
  .DateLastModified
If DateDiff("n", dFileModDate, Now) >
  240 Then
   objShell.run "E:\Jobs\LowSpaceEmail
     .bat"
End If
```

Listing 3: Email_Log_File.bat

```
blat.exe G:\Scripting\Reports\
  LowDriveSpace.log
   -to myemailaddress -subject
     "LowDriveSpace.log has been Updated"
```

Tell the IT community about the free tools you use, your solutions to problems, or the discoveries you've made. Email your contributions to r2r@windowsitpro.com.

If we print your submission, you'll get $100.

Submissions and listings are available online at www.windowsitpro.com. Enter the InstantDoc ID in the InstantDoc ID search box.

An Effortless Way to Confirm the Presence or Absence of Computer Objects

Did you ever have to hunt for a computer object across every domain controller (DC) in your organization to be absolutely sure it no longer existed? Chasing down a computer object can sometimes seem like playing a game of "Where's Waldo?" I tried to find a built-in Windows tool to easily do this, but to no avail. So, I created the Find Computer Object tool.

Here's how the tool came about. Our Desktop Support Team had experienced

some trouble rejoining a computer to our domain. The support team members tried the standard procedures:

- Using the Computer Name tab on the System Properties dialog box, they moved the computer object to a workgroup, then moved it back into the domain, using the administrative credentials required for joining a computer to the domain.
- In the Microsoft Management Console (MMC) Active Directory Users and Computers snap-in, they located and deleted the computer object from the domain, waited 90 minutes, then rejoined it to the domain.

After each procedure, the support team members and the computer's users thought that the problem had been solved because the users could log on to the domain. However, when the users tried to log on later, they found they couldn't. When the support team members checked, they found that the computer object had "dropped off" the domain.

This problem was frustrating both the users and the support team members, so the troubleshooting team that I belong to was brought in to dig a bit deeper. We had the support team once again join the computer to the domain. We watched the computer object appear on the DC. Later, we confirmed that it no longer existed on the same DC.

Harry Verge

A quick ping sweep showed that all the DCs were responding, so we used the Active Directory Users and Computers snap-in to manually search each DC to confirm that the computer object no longer existed on it. That's when we discovered that the computer object still lingered on one DC.

It turned out that this particular DC was having a replication issue with its partners, but for some unknown reason, the monitoring system didn't pick up on the problem. The DC had a corrupt database but was still being advertised to clients as valid through DNS.

Having to manually search each DC to find the problem made me realize that there had to be an easier way to check for the presence or absence of a particular computer object across all DCs. That's when I decided to create the Find Computer Object tool.

The Find Computer Object tool can serve a dual purpose. You can use it to not only determine whether a computer object is present on a DC but also pinpoint exactly where that object resides in the organizational unit (OU) structure if it's present.

Figure 1 shows the tool's UI. After you enter the name of the computer for which you're searching and click the Find Computer button, the tool builds a list of the computer objects in the Domain Controllers OU in AD. The tool then queries each DC on that list, looking for the computer name you entered. It also increments a progress bar so that you know it's actually doing something.

When the tool finds the computer object, it stops and notifies you. As Figure 2 shows, it tells you the DC on which the computer object was found and the OU in which the computer object resides. If it doesn't find the computer object on any of the DCs, you'll receive a message similar to the one shown in Figure 3.

You can download the Find Computer Object tool, which is an HTML Application (HTA), from the *Windows IT Pro* website. Go to www.windowsitpro.com, enter 129545 in the InstantDoc ID box, click Go, then click the *Download the Code Here* button.

—Harry Verge,
senior technology specialist

InstantDoc ID 129545

Figure 2: Message reporting that the computer object was found

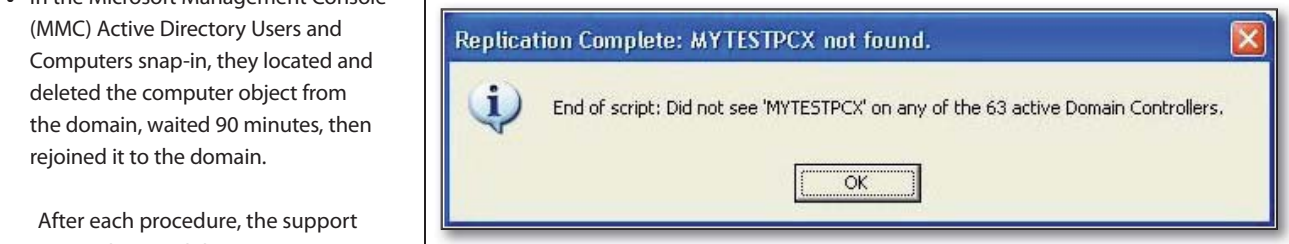

Figure 3: Message reporting that the computer object wasn't found

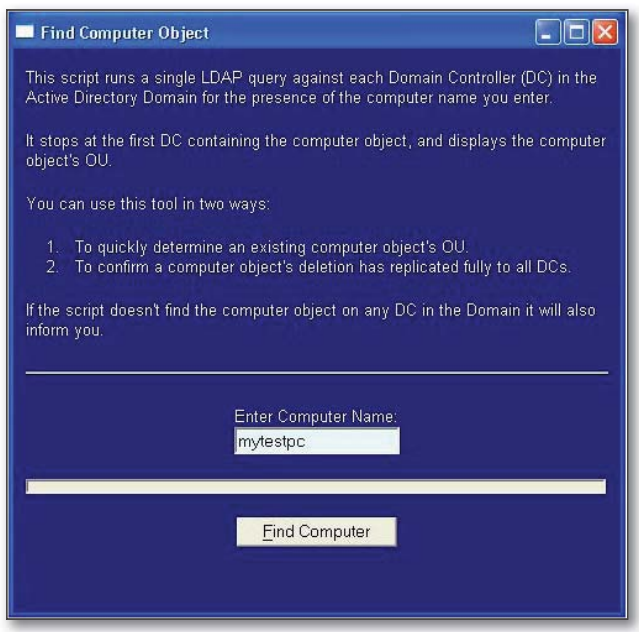

Figure 1: The Find Computer Object tool's UI

The Conversation Begins **Here**

APRIL 17-20, 2011
BELLAGIO, LAS VEGAS

mobile connections | cloud connections | virtualization connections

Questions Answered • Strategies Defined • Relationships Built

NOWHERE ELSE WILL YOU FIND THIS MANY
INDUSTRY EXPERTS

Jay Freeman
CYDIA

Tyler Lassard
RIM

John Stetic
NOVELL

Joe Belfiore
MICROSOFT

Jim Reavis
CLOUD SECURITY ALLIANCE

Emerging trends in cloud computing

How to build, market, and deliver apps seamlessly

How current and future virtualization products will help shape the future of cloud computing

NETWORK WITH YOUR PEERS!

Network with your peers, carriers and a wide range of mobile infrastructure, product and service vendors! Get bet-the-business market knowledge for business leaders, developers and IT managers exploring or implementing cross-platform mobile development, cloud computing and virtualization.

50+ SPEAKERS!
8 KEYNOTE PRESENTATIONS
60+ BREAKOUT SESSIONS
12 BOOT CAMPS

Brought to you by: connected planet | DEVCONNECTIONS | SQL SERVER | Windows IT Pro

3 CONFERENCES - 1-STOP EXPO - GREAT NETWORKING

REGISTER FOR 1 EVENT, GAIN ACCESS TO ALL 3!

Ric Telford
IBM CLOUD SERVICES

Robert Scoble
RACKSPACE

Jinesh Varia
AMAZON WEB SERVICES

Nils Puhlmann
ZYNGA

Ilja Laurs
GETJAR

Getting optimal business efficiency using Microsoft and VMware virtualization solutions

Description of mobile panel: Which platform do you bet on?

ROI of implementing cloud and virtualization platforms and solutions

Aaron Hillegass
BIG NERD RANCH

Michele Leroux Bustamante
IDESIGN INC.

Paul Thurrott
WINDOWS IT PRO

Steve Riley
RIVERBED TECHNOLOGY

Pamela Dingle
PING IDENTITY

Ryan O'Hara
MICROSOFT

REGISTER TODAY!

TheConversationBeginsHere.com or call 800.505.1201

SCSM
BranchCache
Outlook
SSL
XenApp

ANSWERS TO YOUR QUESTIONS

Q: What does Conversation Clean Up in Outlook 2010 actually do?

A: Microsoft Outlook has long allowed users to arrange email by conversation. Outlook 2010, however, improves on that presentation. As users reply to conversations, original content from the thread gets included in responses. The result, especially in a long conversation, is a lot of repeated content across all the messages. Outlook 2010 adds a new feature called Conversation Clean Up. This feature works for any type of account in Outlook 2010; however, if you're using Microsoft Exchange Server, whichever version, Outlook 2010 must be in cached mode.

Conversation Clean Up scans the thread for messages that are wholly contained within other messages, then deletes the older one. To me, this sounds scary. I'm trusting Outlook to remove content based on its assessment that the content is redundant. Microsoft included some options to limit what gets removed. These options are found in the Mail section of the Outlook Options window, which is found by navigating to File, Options. Figure 1 shows the different options for the Clean Up process. Outlook by default prevents Clean Up from removing categorized or flagged messages, messages altered by a reply, and digitally signed messages. You can also opt to exempt messages that haven't been read.

To clean up a conversation, right-click a conversation within any mail folder and select Clean Up Conversation from the context menu. You'll then see a confirmation dialog box. Running Clean Up moves redundant messages to the Deleted Items folder unless you've defined an alternate location, which you can do in the Mail section of Outlook Options. If you don't see the Clean Up option, it's likely because you aren't right-clicking at the top of a conversation. The Clean Up option isn't visible when you right-click a message within the conversation view.

You'll see the full power of this feature at work when you select Clean Up from the Home tab on the Ribbon. From here, you can apply the Clean Up to a conversation, a folder, or a folder and its subfolders. If this option has never been used, or even not recently used, this process can significantly reduce the size of a mailbox. If you're concerned about what's being removed, you can certainly review the Deleted Items folder or whatever custom destination folder you've assigned.

As the Conversation Clean Up feature deletes messages in a thread, it also removes some of the ability to troubleshoot based on those messages on the client side, including working with header information. The body and envelope of a message deleted by the Clean Up process still reside in an existing message; however, the SMTP header information is no longer available. The trade-off is that Conversation Clean Up can help reduce the volume of messages in your mail folders.

—William Lefkovics

InstantDoc ID 129549

Jan De Clercq | jan.declercq@hp.com
William Lefkovics | william@mojavemediagroup.com
John Savill | jsavill@windowsitpro.com
Greg Shields | virtualgreg@concentratedtech.com

Q: What's the correct order to update my System Center Service Manager (SCSM) 2010 installation to SP1 (or any other service pack)?

A: The correct order for upgrading SCSM 2010 is:

1. Upgrade the Service Manager Data Warehouse server.
2. Upgrade the initial Service Manager Management server (the first management server installed).
3. Upgrade the Service Manager consoles, any other management servers, and any self service portals.

—John Savill

InstantDoc ID 129512

Q: How can I disable the Encrypting File System (EFS) to block users from using EFS to encrypt files that are stored on the NTFS file system in Windows 7 or Windows Vista?

A: You can disable EFS using a registry hack. You must modify the value of the NtfsDisableEncryption registry key, which is located in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\FileSystem. To disable EFS you must set the key's value to 1. You can accomplish the same thing by using the following from a command line:

```
fsutil behavior set disableencryption 1
```

In both cases, you must restart your computer to apply the change. Remember that EFS is available only in Windows 7 Professional, Ultimate, and Enterprise editions and Windows Vista Business, Ultimate, and Enterprise editions.

—Jan De Clercq

InstantDoc ID 129588

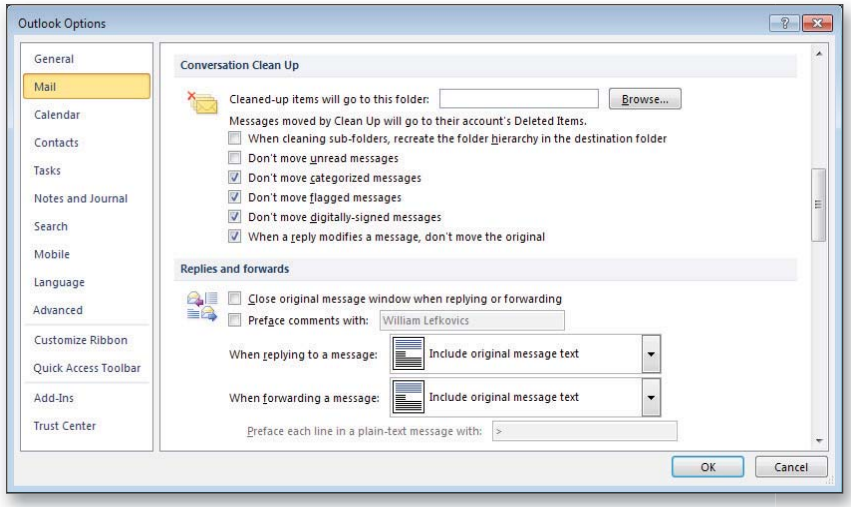

Figure 1: Options you can set in Outlook 2010 for Conversation Clean Up

Q: How can I grant certain users the ability to run System Center Configuration Manager (SCCM) reports but nothing else?

A: SCCM has very granular security capabilities, including the ability to grant users access to only specific reports. First, make sure you have at least one Report Point site server in your environment. On the Report Point site server, there will be a group called SMS Reporting Users. Add users you want to be able to run reports to this group.

I prefer to create a domain global group named SMS Reporting Users and add that domain group into each local SMS Reporting Users group on each Report Point server to simplify management. This way, if additional users need to run reports, you can just add those users to that domain global group.

If users want to run reports from the SCCM console and not the SCCM website, ensure that the SMS Reporting Users domain group also has read access to the Site object class for all instances. If users should access only specific reports, give those users read access to those specific report instances. For example, you can give reporting group access to only two reports. When you go to the website, you'll see only those two reports. Reports are cumulative, so you can give different users different rights, and when they look at reports they'll get the sum of all reports available to them and any groups they're in.

—John Savill

InstantDoc ID 129508

Q: I'm trying to access my smart card in Windows Virtual PC but my VPN software says the card can't be read correctly. How can I make it work?

A: Windows Virtual PC has integration features that allow certain types of devices to be shared between the host computer and the virtual machine (VM), including audio devices, clipboard, printers, drives, and smart cards. However, when you share devices, the device is typically virtualized over RDP. The VM will see a generic smart card and won't function correctly. Because the smart card is shared, you can't directly map the smart card into the VM using the USB attach menu.

The solution is to temporarily turn off the integration features. Attach the USB smart card reader, ensure the right drivers are installed and that it functions, detach the device, then enable the integration features again. The smart card should now function using standard smart card sharing. Below are more detailed steps.

1. Disable integration features.
2. You'll then need to log on again. Now, in the USB menu, attach the smart card device and click Continue when the warning message appears stating that doing so will remove the device from the host computer.
3. Test your VPN or other software after installing the necessary drivers. The smart card should now be functioning correctly.
4. Use the USB menu to release the USB smart card from the VM.

Enable the integration features again, and the smart card should be available and working, because you have the right drivers in the VM.

—John Savill

InstantDoc ID 129555

Q: What's Active Directory (AD) Link Value Replication (LVR) and how does it benefit AD security?

A: One important deficiency of the way that groups are implemented in Windows 2000 AD is that a group's membership attribute is completely replicated between domain controllers (DCs) every time a change in group membership occurs. A change can be as small as adding or removing a single user to or from the group. This is because group membership is implemented as a normal multi-value AD attribute, and multi-value attributes are replicated as a single data blob.

Besides transferring more data over the wire than really necessary, the key problem with this implementation is that when administrators are updating group membership almost simultaneously on different DCs, one administrator's changes will be overwritten by the other administrator's changes as they're replicated between the two DCs. In Win2K AD, the last writer wins and the first writer's changes are lost.

Windows Server 2003 AD replication introduces AD LVR, which resolves this problem. Thanks to LVR, individual values of a multi-value attribute can be replicated separately between AD instances. LVR also reduces AD replication traffic, network bandwidth usage, and processor and memory usage.

LVR is available to you only if your AD forest doesn't include any Win2K DCs. For Server 2003, this means that your forest must be in the Server 2003 interim or native Server 2003 functionality level. For Windows Server 2008, this means that your forest must be in the Server 2003 or Server 2008 functionality levels.

—Jan De Clercq

InstantDoc ID 129487

Q: What's a Worker Group in Citrix XenApp 6?

A: Worker Groups are new in Citrix XenApp 6. They're collections of XenApp servers that reside in the same farm and are managed as a single unit. Worker Groups let you collect servers into groups for publishing applications, load balancing, and policy filtering. They're particularly useful for larger installations where many XenApp servers must be managed as a single unit.

Worker Groups can be created based on each server's Active Directory (AD) organizational unit (OU) membership, or a Worker Group membership can be specifically assigned to each server in the Citrix Delivery Services Console.

By combining XenApp's ability to stream applications to XenApp servers with the configuration control in Citrix policies, Worker Groups make it feasible to add a barebones XenApp server to an OU and see it be automatically configured and provisioned with the Worker Group's applications. Using applications that have already been prepared for streaming delivery, this capability can significantly decrease the amount of time required to deploy new XenApp servers.

—Greg Shields

InstantDoc ID 129494

Q: Is the Windows 7 BranchCache feature used when you use SharePoint 2010 with Office 2010?

A: One key feature Office 2010 and SharePoint 2010 is the use of a new protocol, MS-FSSHTTP, that facilitates synchronization of files via SOAP over the HTTP protocol. This new protocol allows locking of portions of a file, permitting the concurrent co-authoring of the same file. It also lets you download and upload only changes to an Office file, instead of the entire file (which is what occurs when using WebDAV).

A local cache of content previously accessed by Office 2010 from SharePoint 2010 is maintained, allowing only changes to the file to be downloaded on subsequent access. This local cache is managed through the Office 2010 Upload Center application, which lets you modify cache settings as needed. This transfer of only changes to data means very efficient use of network bandwidth.

So where does BranchCache come into this discussion? BranchCache allows the sharing of data that's transferred over HTTP or SMB, so that once data has been downloaded by one computer at a location, it's shared with other clients at the location, reducing data sent down the wire. Because SharePoint 2010 with Office 2010 uses the MS-FSSHTTP protocol instead of normal HTTP and has its own built-in features to handle efficient use of network bandwidth, BranchCache won't be used for the SharePoint 2010/Office 2010 traffic. BranchCache is still used for other types of traffic sent over HTTP and SMB, including data accessed by previous versions of Office.

—John Savill

InstantDoc ID 129558

Q: How does data execution prevention (DEP) work in Outlook 2010?

A: Microsoft introduced DEP to the Windows OS after Microsoft's Trustworthy Computing initiative in 2002. It has become standard since Windows XP SP2. Microsoft has now brought this level of security into Microsoft Office 2010.

In Office 2010, DEP works at the software level to prevent code that doesn't meet requirements from executing. Such code typically originates through Office add-ins. When inappropriate code from an add-in attempts to use memory pages to execute, Outlook 2010 stops working, and appears to have crashed. When you restart Outlook, a warning pops up advising you of a problem with the add-in. The popup also suggests that you disable the specific add-in until it is fixed or updated. Overall, DEP protects your workstation from bad code, whether accidental or malicious, by not allowing the code the memory resources it needs to execute.

DEP can be toggled on or off through the Trust Center in any Office application, including Outlook 2010. In the 32-bit version of Outlook 2010, you access the Trust Center by selecting File, Options, then choosing Trust Center from the sidebar menu on the left. Click the Trust Center Settings button under the Microsoft Outlook Trust Center section. Here, you can select the DEP Settings option and click the check box to disable or enable DEP. (In Office applications that support Protected View, such as Word 2010, the option is found by clicking Protected View in the Trust Center Settings window.) DEP is always enforced in 64-bit versions, so it's not a configurable option. You can change the DEP setting through the registry as well. The DWORD value for the key EnableDEP is 0 for disabled and 1 for enabled. You'll find it in the registry at HKEY_CURRENT_USER\Software\Microsoft\Office\14.0\Outlook\Security.

You can confirm that applications are protected using DEP through the Task Manager. You can view the Data Execution Prevention column in Task Manager by selecting View, Select Columns, and clicking the check box beside Data Execution Prevention.

If your enterprise has an Outlook add-in that you know isn't malicious but keeps causing DEP exception errors, you can disable DEP for Outlook. Ideally, you should get an update for that add-in that doesn't try to execute code from memory pages not intended for code execution.

—William Lefkovics

InstantDoc ID 129612

Q: In Citrix XenApp 6, what's the Citrix Delivery Services Console?

A: Versions of Citrix XenApp before 6 always struggled with needing to spread administrative activities across a series of management consoles. Even as recently as XenApp 5, Citrix hadn't combined all its management functions into a single, unified console—with XenApp 5, some activities were accomplished in the Access Management Console whereas others (most notably, managing policies) were done inside the XenApp Advanced Configuration console.

XenApp 6 unifies all administrative actions in the newly named Citrix Delivery Services Console. This console includes all the functionality of the previous two while adding new functionality such as the ability to manage Citrix policies as Active Directory (AD) Group Policies.

—Greg Shields

InstantDoc ID 129492

Q: In Citrix XenApp 6, when should I use published applications instead of streamed applications?

A: Using Citrix, you have a range of options for how applications are delivered to users. Among the many available options, making the decision between two in particular—published applications and streamed applications—can be difficult to the newbie XenApp administrator.

A published application has been installed on a Citrix XenApp server. Published applications run on the server, consume server resources, and only transfer screen updates and keyboard and mouse commands from the server to the client. Applications that must keep their execution in the datacenter or that have light resource needs are usually good candidates for being published.

On the other hand, some applications are needed when users can't connect to XenApp servers in the datacenter. Other applications use significant amounts of resources when they're used, making server installation a bad idea. When one big application consumes lots of resources, there simply aren't many left to share among other users. These applications make good candidates for streaming to user desktops, because streamed applications consume desktop resources and not server resources as they're used.

Things get less clear when applications are streamed directly to XenApp servers instead of to user desktops. With the features available in today's XenApp software, this can be a really useful practice. Remember that streamed applications install automatically without requiring extra effort at the console. Streamed applications can also be quickly removed, essentially just by deleting a few files. By streaming applications to your XenApp servers, you can easily and automatically deployed as they're needed. Servers are also much more easily reconfigured to meet user demands.

In XenApp, these applications are published as "Streamed to a server." Once streamed, they're then published from that server as published applications.

—Greg Shields

InstantDoc ID 129490

Q: I need a self-signed certificate for my Windows Server 2008 R2 IIS SSL. How do I create it?

A: If you want to enable SSL traffic to your IIS box for test purposes and don't require a certificate to be trusted by external parties, you can create a self-signed certificate. However, clients who access the server will receive a warning that the certificate isn't trusted.

1. Open the Internet Information Services (IIS) Manager.
2. Select your IIS server.
3. In the main IIS section, select Server Certificates.
4. Under Actions, select Create Self-Signed Certificate.
5. Give a name for the certificate and click OK.

The certificate will now be available for use.

—John Savill

InstantDoc ID 129563

Q: I'm trying to install the System Center Services Manager (SCSM) Web Portal on my SCSM Management Server after upgrading to SP1. The installation fails, saying the portal is already installed, but it's not. What can I do?

A: If you originally installed the SCSM 2010 RTM media and then upgraded to SP1, you can't later install the SCSM SP1 self-service portal on the existing box. If you just have SCSM 2010 RTM, you can install the self-service portal on top of it, but once SP1 has been applied, you can no longer add the portal component.

The solution is to install the SCSM self-service portal on a separate server that has IIS installed. You need the IIS role with IIS 6 metabase compatibility, ASP.NET 2.0, basic and Windows authentication, and the .NET Framework 3.5.1 feature enabled on the box. You also need the authorization hotfix installed (as outlined in "Users and applications cannot access authorization rules that are stored in Authorization Manager" (support.microsoft.com/kb/975332).

You can now launch the SCSM 2010 SP1 setup.exe file and select it to install the web portal component. You need to specify the SQL Server Service Manager database instance and database name along with an account to use for connectivity, which will be the Service Manager Services domain account you specified during the Service Manager management server installation. This account must be a member of the local admin group on the SQL Server system hosting the Service Manager database and have the sysadmin database role on the SQL Server system.

The best practice architecture is to separate the self-service portal service on a separate server from the management server, so this isn't actually a bad workaround to not being able to install the portal onto the management server.

—John Savill

InstantDoc ID 129559

Q: What's Citrix Dazzle?

A: Before Citrix Dazzle, administrators were responsible for managing which applications and content were available to which users. While this management works well in smaller environments, it can grow unwieldy as the size and complexity of a Citrix deployment grows large. At some point, the matrix of which users should be assigned which applications grows difficult to manage for IT alone.

Citrix created Dazzle mainly to resolve this problem. While still letting administrators choose which applications are exposed to which users, Dazzle creates a self-service storefront where users can select the applications they need. Dazzle removes some of the effort in managing users and applications. Selecting applications (whether published or streamed) automatically makes them available in the Start Menu or on the desktop.

In addition to giving more power to users to determine which applications they need, Dazzle can also (in combination with other Citrix components) manage licenses and installations to help you identify whether applications are being used appropriately. Ensuring licensing compliance is also handled, to make sure users don't consume more licenses than you currently possess.

—Greg Shields

InstantDoc ID 129491

Ease Cloud Security Concerns with Federated Identity

Securely extend Active Directory's reach into the cloud

by Sean Deuby

For the first time in a long time, the enterprise identity landscape is evolving at its most basic level. There's a new kid on the block, and its name is federated identity. Although the seeds of this change have been around for a while, we just didn't recognize its importance. Federated identity is here to stay, and IT professionals and developers need to learn about it and how it will affect their work in the future.

Why We Need Federated Identity

To understand the growing popularity of federated identity, it helps to look at the challenges that IT professionals and developers face when using traditional identity authentication in the modern IT environment—in particular, the Kerberos protocol. The point behind an identity provider, such as Active Directory (AD), is to centralize identity information for resources to consume. Although identity-oriented IT pros tend to lose sight of it, the purpose of the authentication process is to determine and validate the user's identity in order to gain access to resources.

The Kerberos security protocol (and therefore the AD domains and forests built on it) was designed to work in a fairly secure environment, such as a corporate intranet. The Kerberos protocol, as implemented in AD, provides two components: confirmation of identity and security group membership. If a resource (e.g., a DFS namespace) requires more information, such as site information, it needs to extract that information from another location—AD itself.

However, scenarios that don't require any modification of AD to store more information are pretty simplistic in real life. Microsoft Exchange Server, for example, requires more information about a user than the base AD schema provides. So, AD admins must extend the schema to allow Exchange to store added identity data about its users. Schema extensions aren't done casually; they take time to prepare for and schedule. As a result, other applications might choose to store identity information in databases such as SQL Server or Active Directory Lightweight Directory Services (AD LDS) that don't require the amount of preparation a schema change does.

But what if the users and resources are in two different enterprises—for example, a joint venture or collaboration, or for a Software as a Service (SaaS) cloud application? Do you create and manage the external users' identities by creating shadow accounts in AD, or do your developers create a separate account database to hold them? How do you keep up with the accurate provisioning and deprovisioning of these accounts? What about providing adequate security for these identities against hackers?

Most companies don't want to manage external users' identities and the headaches that go along with that management. If an application is intended to support multiple access scenarios,

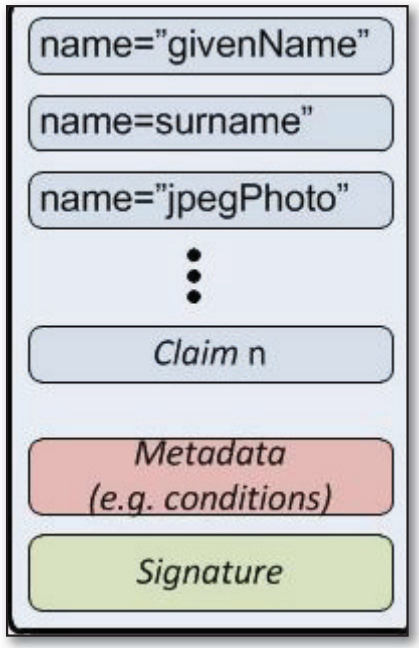

Figure 1: An SAML token

developers must build in multiple authentication mechanisms. Identity design and management in these and other scenarios become very cumbersome, and the traditional model is stretched to its limit.

What Federated Identity Is

The federated identity model can handle a variety of scenarios. Federated identity is the ability to port data across security domains using claims and assertions from a digitally signed identity provider. To understand what that definition means, let's break it into parts. As I described in the previous section, each enterprise's identity store can be generically described as a *security domain,* regardless of whether it's using AD or some other directory product. For the purpose of this article, AD is the *identity provider* for scenarios inside an enterprise. For scenarios that span multiple enterprises, the identity provider is the entire enterprise that provides identity information (not just AD). As for *claims* and *assertions,* these are essential parts of what we call claims-based authentication.

Claims-based authentication is the cornerstone of federated identity. At its simplest, claims-based authentication is about presenting an application with the potentially wide variety of identity information it needs, from an identity provider it trusts, in a highly secure envelope, regardless of whether the application is inside or outside the enterprise. That's why it can handle the two-enterprise and SaaS scenarios that I discussed in the previous section so well. Claims-based authentication adds flexibility and security, whereas traditional authentication technology gives you either flexibility (LDAP queries to AD) or security (Kerberos).

The claims-based authentication model is based on a few simple, intuitive concepts, but the authentication process can bounce back and forth quite a bit. Let's compare some of the basics of this model with one you know a little better: the Kerberos protocol.

In AD, every authenticated user has one or more Kerberos tickets that contain identity information. A basic construct of claims-based authentication is the token, formatted in Security Assertion Markup Language (SAML). Figure 1 shows an SAML token, which is similar to a Kerberos ticket in many ways. A Kerberos ticket contains a payload, called the access token, that asserts what security groups the user is a member of. The resource (e.g., a file server) trusts this assertion because the ticket is cryptographically confirmed to be from a valid identity source—which in AD is the Kerberos Key Distribution Center (KDC) of the domain controller (DC) the file server is talking to.

An SAML token is in fact called an assertion. The payload of this assertion contains a potentially far broader set of identity information, called claims, than a Kerberos ticket holds. An SAML token is designed to be much more flexible in this regard. A claim can be anything you define it to be: name, email, phone number, age, privilege level, meal preference, etc.

In AD, a Kerberos ticket has time restrictions regarding when it can be used. This prevents replay attacks, in which packets are captured then played back to a server at a later time in an attempt to compromise it. An SAML assertion also contains conditions that place more stringent restrictions on when the assertion is valid than the Kerberos protocol is capable of doing. You can restrict when the assertion is valid, who can use the assertion, how many times it can be used, and whether the assertion can be delegated. Thus, an assertion can be highly targeted toward a specific use, and that use only, to increase the security of the authentication process.

Finally, an AD Kerberos ticket is encrypted with either the ticket-granting server key (for a ticket-granting ticket—TGT) or the user key (for a session ticket). An SAML assertion is signed and can have various degrees of encryption from the identity provider that created it, from individual components to the entire assertion. The signing ensures that the assertion is indeed from the stated identity provider, and the encryption ensures that the assertion hasn't been tampered with or spied on.

For all these similarities, though, here's the most important distinction: The scope of an AD Kerberos ticket is essentially within the enterprise, whereas an SAML token has no restrictions of this kind at all. This means that a claims-aware application can authenticate users equally comfortably inside or outside the corporate firewall.

This token doesn't appear out of thin air. Something has to create it, and AD doesn't know anything about this process. Enter yet another component of the claims world: the Security Token Service. The STS issues tokens on behalf of requests from users. Figure 2 shows how an STS interacts with a user and AD to build a token that can be presented to claims-aware applications. Note that Figure 2 explicitly shows the user's browser. This is because it can be closely involved in the process if the user's OS doesn't have a client that understands the token-passing process.

How Federated Identity Works: Two Scenarios

Now that I've introduced all the players involved in claims-based authentication, let's take a look at how all these components work together to authenticate a user to an application. It's kind of a complicated dance, but it's almost entirely transparent to the user. Because of federation's flexibility, federated identity is used in several different scenarios. In this article, I focus on just the two most likely scenarios you'll face. In these scenarios, you must first set up a federated trust between your federation service and the service provider. This necessary step lets the service provider's STS decrypt the encrypted claims coming from your company's federation service; it also lets your STS accept claims requests from the service provider.

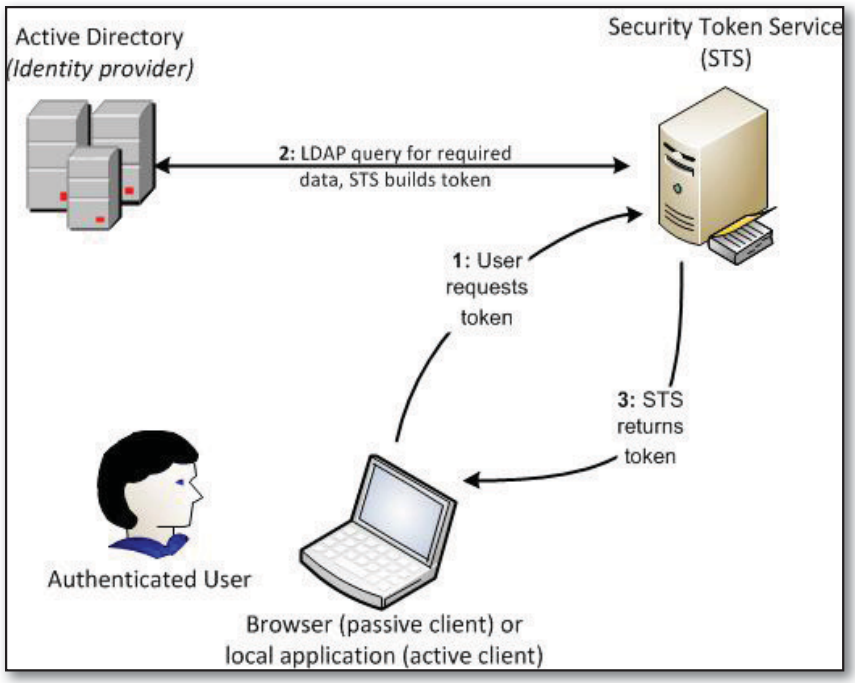

Figure 2: STS token-creation process

The first scenario occurs when a user inside the enterprise attempts to access a claims-aware application that's also inside the enterprise. This situation might not exist in your organization today, but it will be common in the near future as more applications become claims aware and as the private cloud becomes more popular.

The process of an internal user accessing an internal application includes many individual steps, as Figure 3 shows. However, this scenario is easier to understand if you keep in mind the high-level process:

1. The application hits a point at which it can no longer continue (e.g., it needs identity data for the user).
2. The application triggers or initiates either a web service call (if the client is active and has some way of understanding the call) or an HTTP redirect through the browser (if the client is passive and can't handle such a request) to ask for a token from an STS.
3. The STS responds to the request, returning the token to the application.
4. The application is able to continue (e.g., returning data to the user or allowing access to the application).

All you need to implement this scenario is a federation service, such as Active Directory Federation Services (AD FS) 2.0, IBM Tivoli Federated Identity Manager, or Ping Identity's PingFederate, and a claims-aware application such as Microsoft SharePoint 2010.

In the second scenario, which Figure 4 shows, the user is inside the enterprise and needs to access an external web-service provider. There are two major use cases for this scenario. The first use case is accessing an SaaS provider, in which an enterprise uses a service such as Salesforce, Google Apps, or a hosted email provider without maintaining separate passwords at every provider. The second use case is for B2B collaboration, in which users in the identity provider's enterprise need to seamlessly collaborate with users in another enterprise who have documents to share. In this case, the claims-aware application might be SharePoint, which would let users from both enterprises post and work with documents. This overall scenario is broadly known as Internet single sign-on (SSO). Note that in this scenario, the user isn't actively doing anything, and no applications on the local computer are aware of the web service; the user's browser is simply redirecting all the traffic through it. This is what's known as a passive client.

The single largest difference between this scenario and the previous one is that the service provider has its own STS, and the application service trusts it alone. The federated trust agreements that the service provider establishes with its customers are supported by the STS, rather than the application service. This service provider configuration is more scalable than one

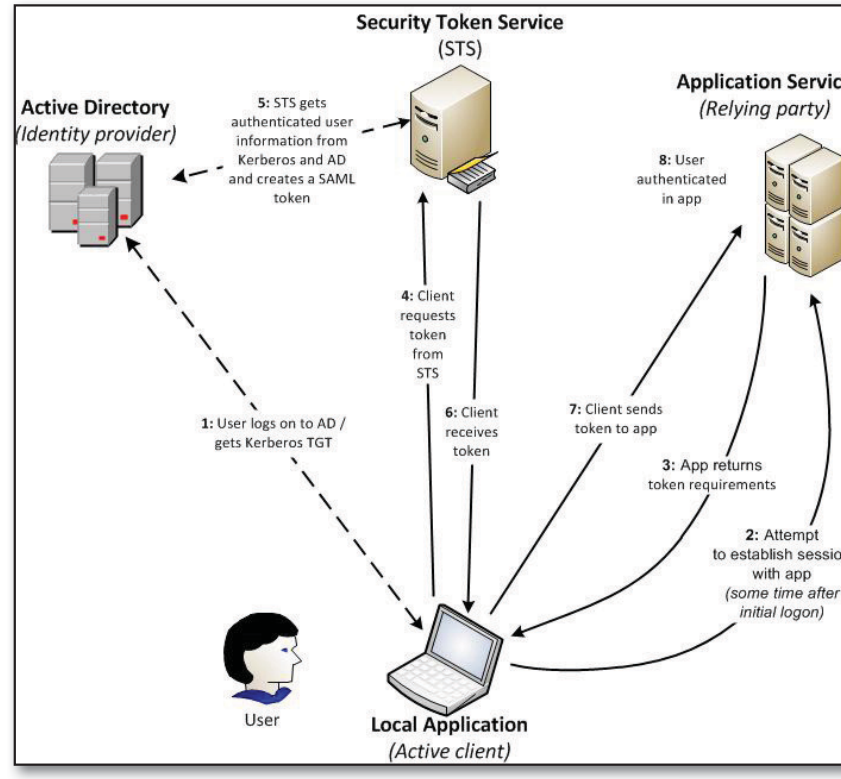

Figure 3: Internal user accessing an internal application

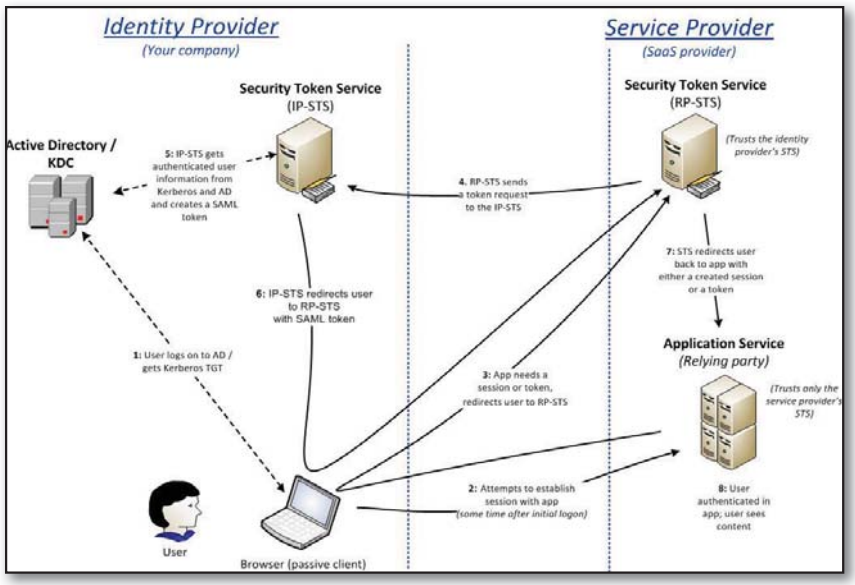

Figure 4: Internal user accessing an external application

without an STS because the resource load of potentially thousands of trusts is focused on the STS instead of the application service and won't affect the application service's resources. It's also more secure because the application service doesn't trust any external claims—only the claims generated by its own STS.

The passive client and the addition of the service provider STS add several steps to the process. Instead of the client actively participating in the authentication process, the application service redirects the request through the client's browser to the service provider's STS to discover the needed claims (step 3). Then, the service provider's STS sends a token request to the identity provider's STS (step 4). After the identity provider's STS generates an SAML token, it redirects the token through the user's browser (neither the user nor the browser has any idea what's going on) to the service provider's STS. This STS will verify it, generate a token with its own signature (the only one the application trusts), and present it to the application service (step 7). The process then completes as expected, and the user is redirected to the application to successfully use it. For more information about claims-based authentication in SharePoint 2010, see *SharePointPro Connections,* "Understanding Claims Based Authentication in SharePoint 2010" (February 2011, InstantDoc ID 128836), as well as Steve Plank's whiteboard video presentation "How ADFS and the Microsoft Federation Gateway work together up in the Office 365 Cloud" at vimeo.com/19177993.

Note that the service provider isn't required to have an STS of its own; the application can directly trust the identity provider. However, this situation might be more common in the B2B collaboration use case, where scalability isn't an issue.

Federation also works if the user in the second scenario is outside the enterprise (e.g., working from home on a nonwork computer without a VPN). Because the user is outside the Kerberos domain, the employer's STS puts up a forms-based authentication page for the user to directly enter enterprise credentials for authentication. After the user is authenticated, the claims-based authentication sequence continues.

How Federated Identity Is Used

Now that you've seen how some of federation's moving parts work, you might wonder whether anyone has gone to the trouble to implement it. The adoption of federation technology was slow in its early years because few companies saw the ROI for internal applications and the occasional external collaboration. It took cloud computing, an increase in the number of claims-aware applications, and the explosion of SaaS providers to really give federation the boost it needed. You're already using it today; you're just not aware of it. (Which, after all, is what federation is all about; if it's doing its job, you shouldn't notice it.) If you use any web services that require Windows Live ID, such as TechNet, MSDN, Windows Live Messenger, or any of the other Windows Live properties, you're already using federated identity in the consumer space.

Many companies are implementing federation to keep pace with their users' demands to use SaaS services in a secure and scalable manner. If you set up a federated trust with the provider, your users can log on to the service using their own user IDs and passwords—they don't have to create and manage a separate account; it's handled automatically. The enterprise's account management team no longer has to worry about managing duplicate accounts for multiple SaaS providers—especially the important security task of deprovisioning accounts that shouldn't be active. And after you set up your federation environment with your own STS, it's a trivial task to add new trusts as you acquire new service providers and applications.

Who are the major vendors in federation and Internet SSO software? Microsoft is certainly one of them. AD FS 2.0 is a free download for Windows Server 2008 R2 or Server 2008 (available at bit.ly/gRQXHI). It does a good job, but it's not a trivial implementation; you'll want to work through the TechNet documentation and step-by-step guides (available at bit.ly/dtaVc1) in a lab environment first. Along with AD FS, Oracle Identity Federation, CA Federation Manager, and Ping Identity's PingFederate comprise the majority of the enterprise Internet SSO market.

There's another class of federation software that sidesteps the need for a local STS installation. Products such as Ping Identity's PingConnect, Symplified, and Okta make federation itself a cloud service. The companies host federation software and manage trusts with a vast number of SaaS vendors so that subscribers to these services automatically have secure access to the vendors.

A fairly small percentage of SaaS vendors accept federation today, but the number is rapidly growing. As federated identity becomes common between enterprises and cloud service vendors, the idea of using claims-based identity for applications inside your company won't seem nearly as radical as it does today. A benefit

of claims-aware applications is that they can coexist peacefully with your existing Kerberos-based infrastructure and applications because the STS translates the Kerberos identity information into claims for the applications. Think of the STS as a proxy or gateway between the Kerberos world and the claims world.

The growth of the market for internal claims-aware applications (enabling traditional applications and creating new ones) is a chicken-and-egg situation. ISVs don't want to invest in making applications claims-aware until there's a good customer demand for such applications. But customers won't generate much demand until they're equipped to support these applications with a federation service and can use this form of authentication with little added expense. What will tip the market into broad adoption is the SaaS scenario that I discussed; companies that add federation capability to support their SaaS vendors are positioned to begin using claims-aware applications internally. To help drive this cycle, you should insist that your SaaS vendors provide federation capability. It reduces the risk for you, increases your visibility into the cloud, reduces the vendor's need to maintain an identity store, and helps to position your business to handle claims-aware applications.

Your Next Step with Federated Identity

The best way for you to get started in understanding federated identity is to start playing with it yourself. Set up a federation service in your lab. (I'll write about my experience installing AD FS 2.0 in my own lab in my June Enterprise Identity column.) Start the project to add a federation service for your company. Approach your company's Information Security team first to gain support; if your security administrators aren't already aware of the risk posed by having separate accounts for every SaaS vendor, you should make them aware. A federation service will lower your company's security exposure by decreasing the number of duplicate accounts with SaaS vendors, decreasing overhead costs if IT is attempting to manage these duplicate accounts, and making life easier for your users with fewer logons to remember. If you don't want to host a federation service yourself, federation-as-a-service products such as PingConnect, Symplified, and Okta will outsource it.

Federated identity is a key enabler to integrating cloud services and on-premises traditional IT services. At the moment, cloud computing's hype outstrips its current use—but don't mistake it for only a fad. Virtualization, the web, and the Internet itself all went through these cycles, and they're an accepted part of our infrastructure today. It's time to begin adding federation skills to your career toolset. Federation's importance will only grow in the future, and these skills will be crucial to both your company and your career.

InstantDoc ID 129610

Sean Deuby
(sean@windowsitpro.com) is technical director for *Windows IT Pro* and *SQL Server Magazine* and former technical lead of Intel's core directory services team. Sean has been a directory services MVP since 2004.

FEATURE

Exchange Server 2010 Role Based Access Control

Increased security, easier management, and more flexible permissioning

by Paul Robichaux

Have you ever heard the expression "It's not what you know, it's who you know"? This is often true in the real world, and it's equally true for computer systems. Most systems rely on the idea that every user has an individual account that is not shared with, or accessible by, other people. That paradigm has been in effect for nearly 50 years, and it still holds up pretty well for most uses. In some cases, however, you don't want to assign permissions or user rights to individual users. Instead, it might be more effective to assign those permissions to someone who holds a particular role.

For example, consider the case of a small company that has three employees who handle accounts. Any one of these people has the ability to sign checks, pay bills, and perform other accounting functions. It seems reasonable that you might assign permissions in Active Directory (AD) or even in Microsoft Exchange Server directly to these users. But consider a company that has 3,000 people in its accounting department. Making individual permission assignments to this many users—either individually or in groups—might not be the most efficient way to handle the situation.

As with so many other problems in computing technology, the solution to this situation has existed for nearly 40 years. Time-sharing computer OSs such as OpenVMS and Multics originated the idea of giving permissions to a role, not a user. By assigning permissions for certain objects to the accounts payable role, for example, and then giving users that specific role, it becomes simpler to understand, monitor, and audit which users have which permissions. Unfortunately, role-based access control (RBAC) was largely ignored as part of the transition from large-scale mainframe systems to commodity servers that run Windows or Linux. Now, Exchange Server 2010 has brought back RBAC with a vengeance.

Exchange and RBAC

Exchange 2000 Server introduced the concept of Exchange-specific roles. The Exchange Administrator, Exchange Full Administrator, and Exchange View Only Administrator roles could be used to assign capabilities to individual users or groups. These roles were a fine idea, but they were limited because of the way in which they were implemented. There were two primary limitations. First, the roles were not very granular, making it difficult to assign the exactly correct access or to assign many users. Second, the Exchange management tools themselves had to rely on Windows for authentication information. Windows itself doesn't implement all the underpinnings that are required for RBAC. Accordingly, Exchange couldn't do much more than impose some access checks to verify whether a user account held one of the administrator roles.

Exchange 2010 changes this completely. As in Exchange Server 2007, every action that a user or administrator can take in Exchange Management Shell (EMS), Exchange Management Console (EMC), or Exchange Control Panel (ECP) actually calls one or more PowerShell cmdlets. Exchange

Role Based Access Control (RBAC) lets administrators control which specific PowerShell commands—and even which arguments to commands—an individual user can execute. This is quite a departure from the old model, so it will require some explanation.

RBAC Basics

RBAC depends on three sets of definitions: management role, management role group, and scope. These can be described as follows.

A management role specifies what can be done. For example, Exchange 2010 defines roles for unified messaging (UM) management, discovery management, and other administrative operations. Users who hold one of these roles can take specific actions that are allowed in the definition of that role. In Exchange, the role is made up of role entries, each of which defines an EMS cmdlet or set of parameters that users who hold that role can execute. Exchange 2010 SP1 includes definitions for about 70 roles: for moving mailboxes, working with mail recipients, managing legal holds, and so on. Some of these roles are intended for administrators, but others are intended directly for users. For example, the MyBaseOptions role is typically assigned to users so that they can edit some of their own contact details by using ECP. To get a list of Exchange roles, you can run the Get-ManagementRole cmdlet from an Exchange 2010 server.

A management role group specifies who can do something. This nomenclature is a little confusing. You might think that membership is defined in the role itself, but the role group defines membership. Role groups are actually Windows universal security groups that are stored in the Exchange security group's organizational unit (OU) in AD. You can see the group membership by using a tool such as the Microsoft Management Console (MMC) Active Directory Users and Computers snap-in. However, you should not edit the group membership manually. Group memberships contain custom attributes that would likely be damaged if you were to edit them by using the Active Directory Users and Computers snap-in.

A scope defines a set of objects on which a role can take action. For example, you might define the scope of a given role to be an individual mailbox database, an Exchange server, or an OU. Scopes can grant read or write access to a set of objects.

The role assignment ties these definitions together. When you assign a user to a particular role, the user gains the ability to execute the commands that are defined in that role against a particular scope. For example, you might grant an administrator the organization management role to allow a user to execute a wide variety of commands on a wide variety of Exchange objects throughout the Exchange organization. You might, then, assign a different user a more limited role, such as the ability to conduct organization-wide mailbox searches.

It's important to remember that RBAC assignments are applied by using their own mechanism, not by using the standard Windows ACL mechanism. Usually, if you've defined multiple sets of permissions on a resource, the most restrictive set of permissions is applied. But when you apply RBAC, a user gets the union of all the RBAC role entries to which the user has access. For example, if you assign Joe User two different RBAC roles, Joe User will be able to use the cmdlets that are specified in either of the defined roles, not only those that belong to the most restrictive set.

Figure 1 illustrates the relationship between these components. Microsoft refers to this arrangement as the Triangle of Power—an apt name, considering the way in which the elements are related.

The scope, role, and role group assignments are linked by the role assignment itself. Microsoft recommends that you first define the scope, then the role, then the role group, and, finally, the role assignment. That's the pattern that I'll follow here.

RBAC Scopes

Role scopes define where a role can be applied. Exchange itself defines a wide range of scopes. To help you grasp the full range, I'll first draw a distinction between read and write scopes. The difference is easy to understand: If an object is within the read scope of a role, role holders can read the object. The concept is the same for the write scope.

Each role has separate read and write scopes for recipients and for configuration. That is, every role has a recipient read scope, a recipient write scope, a configuration read scope, and a configuration write scope. The built-in Exchange management roles have implicit scopes. For example, the Public Folders management role has an implicit recipient read and write scope of Organization and an implicit configuration read and write scope of OrganizationConfig. That tells us that someone who holds that role can read and modify both recipient and configuration objects anywhere in the organization. (Remember: The role itself can restrict which specific cmdlets and parameters can be used.)

You can also set your own explicit scopes, either by using a predefined set of scopes that are built into Exchange or by defining your own. The predefined scopes are as follows and can be used only for setting the recipient read and write scopes:

- Organization (global to the organization)
- Self (allows access only to the current user's mailbox)
- MyDistributionGroups (allows access only to distribution groups owned by the current user)

Custom scopes are much more interesting. Exchange supports three different types of custom scopes:

- OU (allows recipient objects in a given OU to be modified)
- Recipient filter (lets you specify a recipient filter that selects the objects

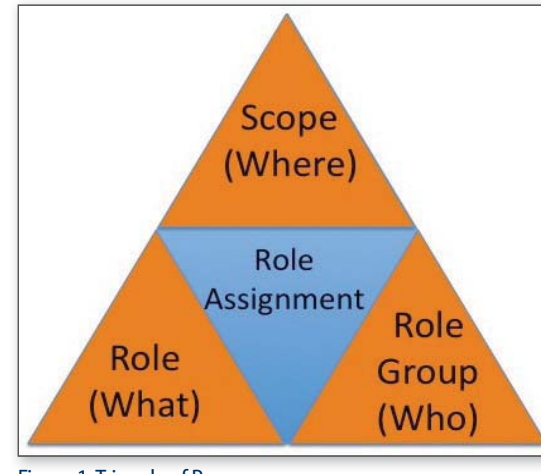

Figure 1: Triangle of Power

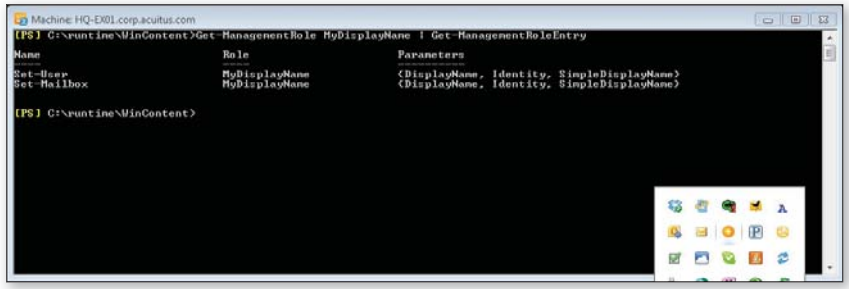

Figure 2: Entries assigned to the MyDisplayName role

that the role holder is allowed to see or change)

- Configuration (allows you to grant read or write access to a subset of objects from the configuration container)

Configuration scopes are the most complex of the three. They allow you, for example, to define a mailbox database scope that restricts role holders to a subset of mailbox databases in an organization. Configuration scopes can be based on filters that dynamically select the objects that should be included, or they can be based on lists, in which you specify the exact set of objects that should be included. Perhaps an example will help clarify the difference. The following cmdlet defines a configuration scope based on a database filter:

```
New-ManagementScope -Name "HQ
  Databases" -DatabaseRestrictionFilter
  {Name -eq "HQ*" }
```

This scope will automatically include all databases in the organization whose names begin with HQ.

Suppose that you then wanted to define a scope that would apply only to three specific hub servers in your HQ site. You could use a list scope to do this:

```
New-ManagementScope -Name "HQ Hub
  Servers" -ServerList HQ-HT01,HQ-
  HT02,HQ-HT-SPARE
```

You've probably figured out by now that you create scopes by using the New-ManagementScope cmdlet. You can get a list of the defined scopes by using the Get-ManagementScope cmdlet, and you remove scopes that aren't assigned to any role assignments by using the Remove-ManagementScope cmdlet.

Roles

After you've defined the scopes you want to use, the next step is to define the roles themselves. Each role defines a set of cmdlets and parameters that role holders are allowed to use. Let's take a look at a simple role.

The MyDisplayName role allows users to change their own display names through ECP. If you use the Get-ManagementRoleEntry cmdlet, you can see the specific entries that are assigned to this role, as Figure 2 shows.

This role contains two entries: one for Set-User and one for Set-Mailbox. Why are there two? Because both of these cmdlets can be used to change the user's display name. In general, you'll see multiple similar role entries for any operation that can be accomplished in more than one way. Many cmdlets (Set-OWAVirtualDirectory, for one, comes to mind) are specific to a single function, so you won't typically see overlapping entries for them.

What about a more complex role? The Unified Messaging role is a decent example. Figure 3 shows part of what the output from Get-ManagementRoleEntry "Unified Messaging*" looks like. There are a few noteworthy things about this management role. First, you can see that there are entries for each of the UM-related cmdlets. Holders of this role can execute any cmdlet that's defined in a role entry within the role, so it makes sense that all the UM-related cmdlets would have entries here so that role holders can use them. Additionally, each role entry specifies a set of parameters to which the user should have access.

It's probably more interesting to dig into a single management role entry to see how the parameters are specified for a given cmdlet. Take a look at the output from the following cmdlet:

```
Get-ManagementRoleEntry "Unified
  Messaging\Set-UMDialPlan" | select
  parameters | fl
```

Figure 4 shows the output. Notice that some of these parameters are actually arguments to the cmdlet, whereas others (such as Confirm, DomainController, and Debug) are common PowerShell switches that aren't specific to this individual cmdlet. One of the beauties of RBAC is that you can control access down to the individual parameters passed to the cmdlet. We'll see shortly how to actually do that.

EMS interprets the role entries when it's asked to initialize for a particular user. It loads cmdlets that have a role entry defined for one of the roles that the user holds. If a given cmdlet doesn't have any matching role entries, it isn't loaded. EMC

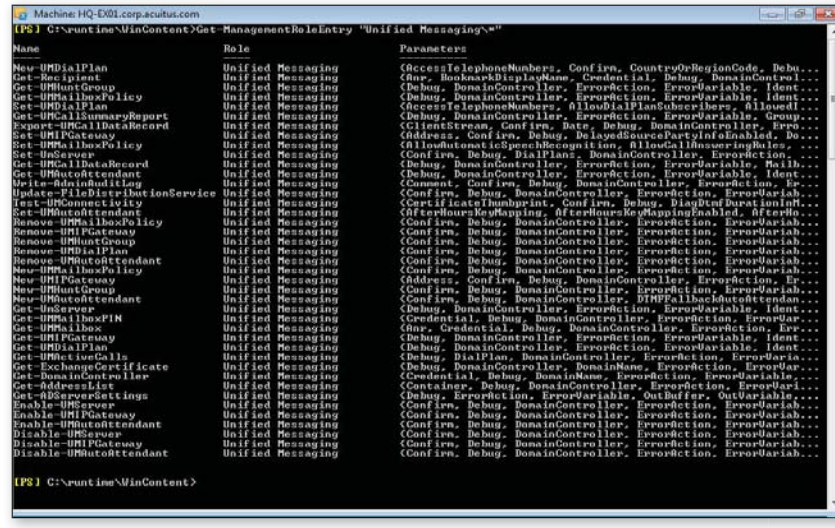

Figure 3: Partial output from Get-ManagementRoleEntry "Unified Messaging*"

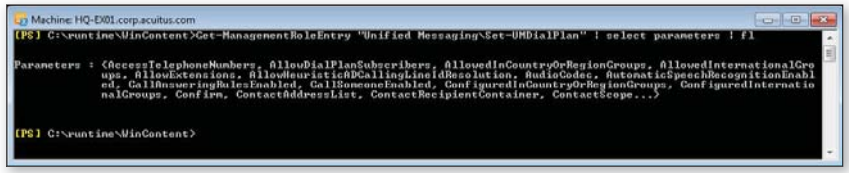

Figure 4: Output from Get-ManagementRoleEntry "Unified Messaging\Set-UMDialPlan" | select parameters | fl

and ECP hide or disable the user interface for cmdlets that are found to be missing during initialization.

With such a variety of built-in roles, the odds are good that you'll be able to find a role that meets your needs. If not, you can create a new role that's based on one of the existing roles, then customize the role by adding or removing cmdlets and parameters.

You can create new custom roles only via inheritance, not from scratch. Because of this restriction, a child role can never have more rights than its parent, though you are free to remove things from the parent. For example, let's say that you wanted to define a custom role that would allow holders to change only the UM PIN for selected mailboxes. The simplest way to do this would be to create a new custom role based on the existing UM Mailboxes role, then remove the cmdlets that you didn't want.

To do this, first create the new custom role by using the New-ManagementRole cmdlet, as follows:

```
New-ManagementRole -name "UM PIN
  Reset" -parent "UM Mailboxes"
```

After you create the new custom role, if you run Get-ManagementRoleEntry on the new role, you'll see that it has access to several cmdlets that you don't want, including Set-UMMailbox, Set-ADServer Settings, and Set-MailboxJunkEmail Configuration. This isn't necessarily bad because each role entry restricts users to a specific set of items. However, if you want to pare things down, you can do so by removing almost all the role entries. You can't remove all of them! If you do, the role loses its ability to do anything and becomes useless. This cmdlet does the trick:

```
Get-ManagementRoleEntry "UM PIN
  Reset\*" | where {$_.Name -ne
  "Get-UMMailboxPIN"} | Remove-
  ManagementRoleEntry
```

This example leaves the Get-UMMailbox PIN role entry intact. It's important to know that you can't have a role entry for a Set cmdlet without also having the corresponding Get cmdlet. In this case, that means that you'll have to restore the Set-UMMailboxPIN role entry, as follows:

```
Add-ManagementRoleEntry "UM PIN Reset\
  Set-UMMailboxPIN" -parameters
  LockedOut, NotifyEmail, Pin,
  SendEmail
```

Role Groups

Role groups themselves are very straightforward. The Get-RoleGroup cmdlet shows you the eleven built-in role groups, which Table 1 describes.

If you follow Microsoft's recommended workflow for using RBAC, the step for defining role groups is easy to follow because you've already defined the "where" (the scope) and the "what" (the role and its role entries) that you want to use. Defining the group lets you specify who is able to work

Table 1: Exchange Server 2010 Built-in Role Groups

Role Group	Description	Related TechNet Article
Delegated Setup	Members of this group can deploy Microsoft Exchange Server without being able to administer it. However, the first server that is running Exchange has to be installed by an admin who is a member of the Organization Management role group.	technet.microsoft.com/en-us/library/dd876881.aspx
Discovery Management	Discovery Management role group members can search mailboxes to retrieve data for discovery requests.	technet.microsoft.com/en-us/library/dd351080.aspx
Help Desk	Help Desk role group members can perform some limited recipient management tasks. For example, they can see and change some parts of the mailbox regional configuration, add and remove Inbox rules, and so on.	technet.microsoft.com/en-us/library/dd876949.aspx
Hygiene Management	Hygiene Management role group members can manage antivirus and anti-spam settings.	technet.microsoft.com/en-us/library/dd776125.aspx
Organization Management	Holders of this role have administrative access to the entire Exchange 2010 organization, so they can do almost anything to any Exchange 2010 object. This is equivalent to the Exchange Full Administrator role in previous versions of Exchange.	technet.microsoft.com/en-us/library/dd335087.aspx
Public Folder Management	Public Folder Management role group members can manage public folders and databases on Exchange 2010 servers.	technet.microsoft.com/en-us/library/dd876947.aspx
Recipient Management	Members of this group get administrative access to create or modify recipients such as contacts and mailboxes.	technet.microsoft.com/en-us/library/dd298028.aspx
Records Management	Records Management group members can configure compliance features, including transport rules, message classifications, and retention policy tags.	technet.microsoft.com/en-us/library/dd633492.aspx
Server Management	Members of this group can manage Exchange 2010 server settings, but they have no access to recipients or other non-server objects.	technet.microsoft.com/en-us/library/dd876866.aspx
UM Management	UM Management members can access and manage UM settings and UM-related recipient properties.	technet.microsoft.com/en-us/library/dd351142.aspx
View-Only Organization Management	Holders of this role can view the properties of any object in the Exchange organization. This is similar to the Exchange View-Only Administrator role from previous versions.	technet.microsoft.com/en-us/library/dd351130.aspx

in the defined scope on the defined target objects.

You can create a new role group by using (no surprise) the New-RoleGroup cmdlet. This cmdlet requires that you specify the group name and one or more roles that you want the group to have. If you want, you can also define a scope for the group. This lets you define a role and assign it to a role group that has a more restrictive scope than the original role definition. However, remember that users in a role group will get access to all the cmdlets and parameters that are allowed under any role in the group. If you don't assign a scope, the default scope for the role is used.

To extend our UM example a bit, you can create a role group for UM PIN management in the following way:

```
New-RoleGroup "UM PIN Managers" -roles
  "UM PIN Reset"
```

You can assign members to the role group by using the members switch:

```
New-RoleGroup "UM PIN Managers"
  -roles "UM PIN Reset" -members
  paulr,brianh,davidw
```

You can also use aliases or SMTP addresses to assign membership.

Role Assignments

Adding the members switch to the New-RoleGroup cmdlet creates management role assignments. This is a useful shortcut. Following the Microsoft model, though, what if you want to create the management role assignments separately?

When you create a management role assignment, you bind a role to a scope and also to a role group or to a user. You can see this for yourself by doing the following: Issue the Get-ManagementRole Assignment cmdlet, use the New-Management Role Assignment cmdlet to create a new assignment, then run Get-ManagementRole Assignment again. You'll notice that a new entry is appended. Exchange 2010 SP1 includes 164 management roles. So when you create new assignments, you see them tacked onto the bottom of the list.

The built-in role assignments all have names that follow a simple pattern: Role-Role Group (e.g., Message Tracking-Records Management). You can use the Get-ManagementRoleAssignment cmdlet to see the details of a specific role assignment by passing its name. The resulting output will show you which roles are assigned, what scopes are in effect for the assignment, and some other useful information about the role assignment.

> RBAC was largely ignored as part of the transition from large-scale mainframe systems. Exchange 2010 has brought back RBAC with a vengeance.

Sometimes, it's easiest to create role assignments when you create a new role group. However, one of the nicest things about Exchange 2010 RBAC is that it includes a very large set of roles and role assignments, so you can almost always use the predefined entries as a starting point.

Role Assignment Policies

Besides the ability to explicitly give users role assignments, you can also give them the appropriate level of access by relying on a little-understood Exchange 2010 feature called role assignment policies. A role assignment policy specifies the role assignments that you want to apply to a group of users. Exchange includes a default role assignment policy that grants users the right to modify most of their own personal information; this default policy applies the MyBaseOptions, MyContactInformation, MyProfileInformation, MyVoiceMail, MyTextMessaging, MyDistributionGroupMembership, and MyDistributionGroups roles to all users.

The beauty of this system is that, by adding or removing roles to this policy, you control the user options in ECP. You can also define new role assignment policies and change the default role assignment policies that are applied to users so that they get exactly the access you want them to have. To do this, create a role assignment policy and use the Set-RoleAssignmentPolicy cmdlet together with the IsDefault switch. Additionally, you assign management role assignment policies to new mailboxes when you create or enable them through the New-Mailbox or Enable-Mailbox cmdlets. You can also use Set-Mailbox to assign a new role assignment policy to an existing mailbox.

RBAC and Exchange Control Panel

So far, I've focused on using EMS to manipulate RBAC objects. In Exchange 2010 SP1,

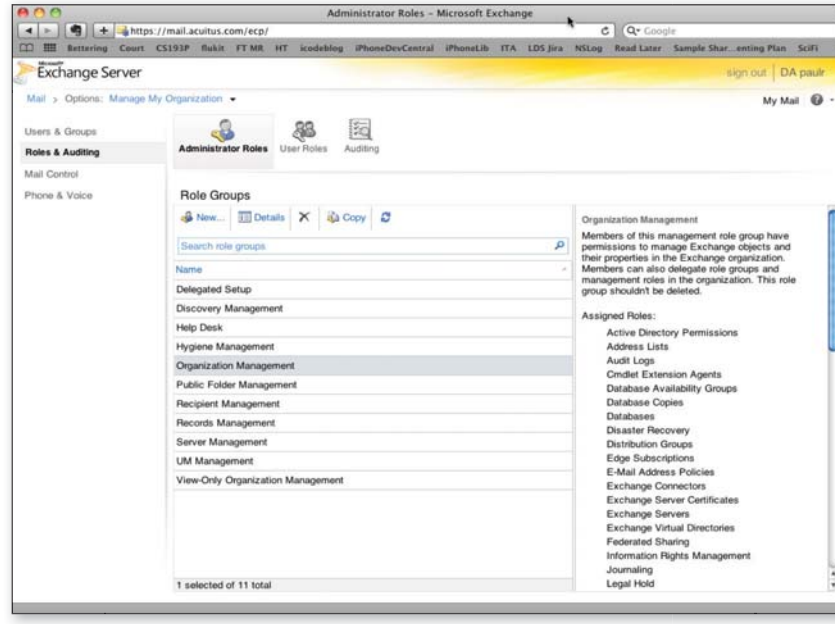

Figure 5: Exchange Control Panel interface for viewing and managing roles and role assignments

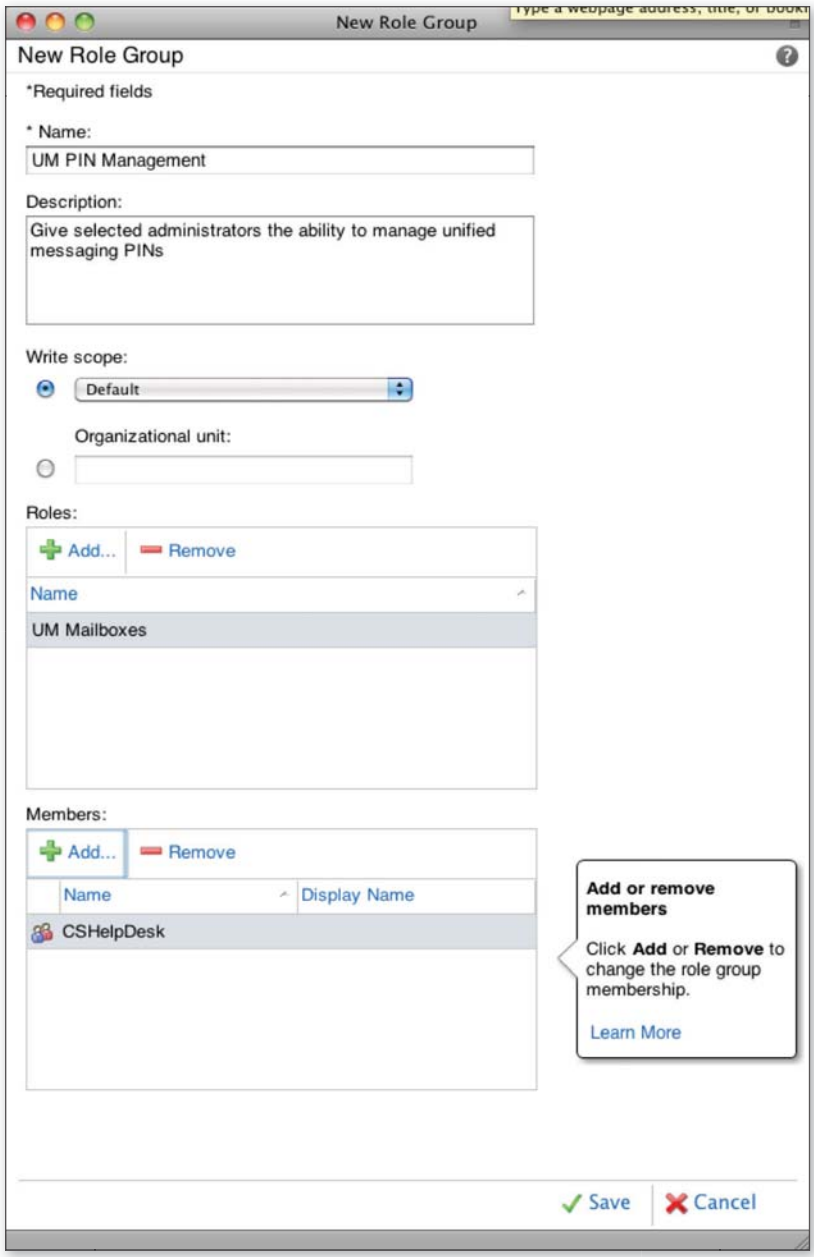

Figure 6: Creating a new role group in Exchange Control Panel

Microsoft added support for some RBAC operations in ECP. This is a great step forward because it makes RBAC much easier to access and to manage for the average administrator.

When you log on to ECP by using an account that has the Organization Management role, you'll see a link labeled Roles & Auditing in the navigation pane. Click this link, and you'll see two unfamiliar icons at the top of the details pane for Administrator Roles and User Roles, as Figure 5 shows. The controls for Administrator Roles let you view, modify, and define role groups for administrative access. The corresponding controls for User Roles give you the same power for user-facing role assignment policies.

Notice that I didn't say anything about roles or scopes here. ECP doesn't include a facility for creating or managing roles or scopes. You can create new role groups by using the existing set of roles and scopes. Also, roles or scopes that you create yourself by using EMS will appear here. That makes ECP a good tool for simple RBAC management tasks, but you can't use it as a 100-percent replacement for EMS.

When you create a new role group, you specify a name, the set of roles that the role group should contain, and the members that you want to include in the role group, as Figure 6 shows. You can't remove cmdlets or change the assigned set of parameters here because you can't change the contents of the roles that you're assigning. Therefore, using this functionality requires some forethought to decide what abilities you want members of the role group to have.

You can use the User Roles tools in ECP to modify the default role assignment policy or to create a new one. Either way, when you modify or create a role assignment policy, you see a dialog box that resembles the one in Web Figure 1 (www.windowsitpro.com, InstantDoc ID 129219). You can control which of the built-in "My" roles users can access. In this example, I've created a policy that lets users manage most aspects of their own account, but not of their distribution group ownership or of their membership. Additionally, some capabilities that I don't want users to have access to (such as the ability to manage text messaging settings) are disabled. This is similar to the default policy that Microsoft uses for its own Exchange operations.

A New Way to Manage

RBAC is a long-established technology that will, nonetheless, be completely new to many Exchange administrators. It seems quite daunting at first compared with other Exchange technologies. This is partly because Exchange 2010 features such as database availability groups (DAGs) build on a foundation of things we already understand, including mailbox databases. RBAC, however, is mostly new.

The time and effort that you invest in learning how to make effective use of RBAC will pay off in increased security, easier management, and a much more flexible permissioning model. I expect that as RBAC becomes better understood, we'll see it become one of the most important drivers for Exchange 2010 adoption.

InstantDoc ID 129219

Paul Robichaux
(probichaux@windowsitpro.com) is a senior contributing editor for *Windows IT Pro*, a content author at Acuitus, and a Microsoft Exchange MVP and MCSE. Paul is the author of *Exchange Server Cookbook* (O'Reilly and Associates), and blogs at www.robichaux.net/blog.

FEATURE

BitLocker Deployment

3 steps to confidence

by Jan De Clercq

BitLocker Drive Encryption (BDE), or BitLocker, offers volume-level data encryption for data stored on Windows clients and servers. BitLocker protects the data when the Windows systems are offline (i.e., when the OS is shut down) and can prevent data breaches such as the theft of confidential data on laptop computers.

In the first version of BitLocker that shipped with Windows Vista, only a single volume, the OS drive, could be protected by BitLocker. Microsoft added support for BitLocker protection of different volumes, including local data volumes, in Windows Server 2008 SP1 and Vista SP1. In Server 2008 R2 and Windows 7, Microsoft added BitLocker support for removable data volumes, memory sticks, and external data drives. Microsoft refers to this feature as BitLocker To Go (BTG).

BitLocker is a great security add-on to the Windows OS; it helps organizations save money because they don't need to invest in special third-party disk encryption software. But organizations are often reluctant to implement new security features, primarily because the features lack a proven track record. Also, new cryptographic solutions bring a certain administrative fear factor to administrators and operators.

To give you more BitLocker confidence, this article highlights three critical steps that you must pay special attention to if you are considering deploying BitLocker in your Windows environment. BitLocker is available in the Ultimate and Enterprise editions of Windows 7 and Vista and in all Server 2008 R2 and Server 2008 editions, with the exception of the Itanium edition.

Choose the Correct Unlock Method

The strength of the protection BitLocker offers depends to a large extent on the authentication mechanism it uses for unlocking access to a BitLocker-protected drive. In BitLocker terminology, this authentication mechanism is referred to as the unlock method.

Before a BitLocker drive is unlocked, BitLocker authenticates the drive based on identification data that the user or the OS provides and that authorizes BitLocker to unlock access to the drive. BitLocker supports different unlock methods based on user knowledge of a secret, presence of a hardware component, or software keys—or a combination of all three. You can select the unlock method when you set up BitLocker.

The available unlock methods differ for OS drives and for fixed or removable data drives. For example, only an OS drive can be protected using a Trusted Platform Module (TPM), a special security chip that is part of most of today's PC motherboards. On an OS drive, you can choose one of the following unlock methods:

Learning Path

More articles about BitLocker:

"Using BitLocker Without a Trusted Platform Module," InstantDoc ID 128895

"Comparing BitLocker with EFS," InstantDoc ID 95603

"Recovering Bitlocker Keys from Active Directory," InstantDoc ID 101582

"A Better BitLocker: BDE Enhancements," InstantDoc ID 102534

"BitLocker and AD, Together at Last," InstantDoc ID 94906

Q&As about BitLocker:

"Q. What is BitLocker To Go?" InstantDoc ID 101445

"Using BitLocker, TPM, and RODCs to Prevent the Exploitation of a DC," InstantDoc ID 98101

"What is the BitLocker exposure via cold boot attack?" InstantDoc ID 99144

"Q. How do I enable BitLocker from the command line?" InstantDoc ID 98219

"Q. Can a USB device encrypted with BitLocker To Go be used on multiple machines?" InstantDoc ID 101446

"Q. Can I read BitLocker protected removable media on Windows XP and Windows Vista machines?" InstantDoc ID 103345

"Q. If I unlock a BitLocker protected USB device, is it only unprotected for the current user?" InstantDoc ID 102857

"Q. How do I configure my BitLocker recovery password to be stored in Active Directory (AD)?" InstantDoc ID 99661

"Q. How can I set the password complexity for removable BitLocker protected devices?" InstantDoc ID 103287

"Q. How can I set the recovery options for removable BitLocker protected devices?" InstantDoc ID 103281

"Q. Does Hyper-V support Windows BitLocker Drive Encryption running on the host?" InstantDoc ID 101746

"Q. How can I stop local administrators from turning off BitLocker?" InstantDoc ID 126076

"Q. Can I use BitLocker on the partition that contains the virtual hard disks (VHDs) from which I natively boot?" InstantDoc ID 125071

"Q. How can I force Windows 7 clients to use BitLocker To Go before writing to USB devices?" InstantDoc ID 103280

- TPM only
- Startup key only
- TPM plus PIN code
- TPM plus startup key
- TPM plus PIN code plus startup key

The last three of these unlock methods offer the best protection. Unlock methods involving a PIN require the user to provide a PIN code at system startup time. When a startup key is involved, at startup time the user must insert a USB token that holds the startup key.

On a fixed or removable data drive, you can choose the following three unlock methods: password, smart card plus PIN, or automatic. For data drives, the smart card plus PIN unlock method offers the strongest protection.

When you use a TPM-based unlock method to protect your OS drive, BitLocker provides integrity checks for critical system files, in addition to data encryption, at boot-up. However, using a TPM adds setup and management complexity and overhead. For example, the TPM must be enabled in BIOS. On most systems, this can only be done after you define a BIOS password. The TPM architecture also requires that an owner password be defined before the TPM can be used. The owner password allows for the clearing and disabling of a TPM and is typically owned by a systems administrator. When you consider deploying BitLocker with a TPM, you must make sure that your computers have a TPM 1.2 chip and a BIOS that is compatible with TPM 1.2 or later. To determine whether a computer includes an operational TPM chip that can be used for BitLocker, check the Microsoft Management Console (MMC) TPM Management snap-in (tpm.msc).

Because many organizations still have older computers that don't have a TPM and you cannot simply add a TPM to a computer, Microsoft included the startup key–only unlock method for OS drives. To use this unlock method, you must make sure that your users have a USB drive and that the computer BIOS supports the reading of USB devices during computer startup. For more information about how to set up BitLocker without a TPM, see "Using BitLocker Without a Trusted Platform Module" (December 2010, InstantDoc ID128895).

When you plan to unlock your BitLocker-protected data drives with a smart card, you must make sure that your users have BitLocker-compatible certificates loaded on a smart card. To generate these certificates, you can use a certification authority (CA), create self-signed certificates, or configure an existing EFS certificate for use with BitLocker. When using smart cards, it is also recommended that you have smart card management software in place. For example, you can use the smart card management functionality that is offered by Microsoft ForeFront Identity Manager (FIM). When you consider using smart cards, I would advise you to carefully read the Microsoft TechNet articles "Using Certificates with BitLocker," at technet.microsoft.com/en-us/library/dd875548(WS.10).aspx, and "Using Smart Cards with BitLocker," at technet.microsoft.com/en-us/library/dd875530(WS.10).aspx.

Create a Solid Recovery Strategy

An encryption tool such as BitLocker requires a solid recovery strategy, and BitLocker forces you to define a recovery method during setup. This will allow you to regain access to the data on an encrypted drive when the drive cannot be accessed (i.e., when the unlock methods that we discussed in the previous section fail).

On an OS drive, you will need a recovery method when a user forgets the PIN or loses the USB token that holds the startup key, or if the TPM registers integrity changes to the system files. For data drives, you will need a recovery method when a user forgets the password or loses the smart card. Also, if a protected data drive is configured for automatic unlocking, you will need a recovery method if the auto-unlock key stored on the computer is accidently lost—for example, after a hard-disk failure or reinstallation. BitLocker supports three recovery methods: a recovery password, a recovery key, and a data recovery agent (DRA).

A recovery password is a 48-bit numerical password that is generated during BitLocker setup. You can save the recovery password to a file, which you then preferably store on a removable drive. You can also print the password, or it can be automatically saved in Active Directory (AD). If you want to automatically store recovery

passwords in AD, you must make sure that all computers can connect to your AD when they enable BitLocker. Storage of BitLocker recovery information in AD is based on an AD schema extension that creates extra attributes to attach BitLocker recovery information to AD computer objects. Server 2008 R2 and Server 2008 domain controllers (DCs) include this extension by default. On Windows Server 2003, you must install the BitLocker-specific schema extension that can be downloaded from www.microsoft.com/downloads/en/details.aspx?FamilyID=3a207915-dfc3-4579-90cd-86ac666f61d4.

To facilitate viewing and retrieving BitLocker recovery passwords from AD, Microsoft provides the MMC Active Directory Users and Computers snap-in. It adds a BitLocker Recovery tab to the properties of the AD computer object. The tab shows all BitLocker recovery passwords associated with a particular computer object. For Server 2008 R2, the BitLocker Active Directory Recovery Password Viewer tool is an optional feature included in Microsoft Remote Server Administration Tools (RSAT). For Server 2008, this extension can be downloaded from www.microsoft.com/downloads/en/details.aspx?FamilyID=2786fde9-5986-4ed6-8fe4-f88e2492a5bd.

The second recovery method uses a 256-bit recovery key that you can save to a USB token or another location. Similar to a recovery password, a recovery key lets users regain access to their protected drive without administrator intervention. When using a recovery key, users must insert a USB token or provide a pointer to another key location during recovery.

The third recovery method, based on a DRA, always requires intervention of a member of the IT department. This method leverages a special certificate that is issued to a dedicated DRA administrator in your organization. The DRA certificate's thumbprint is distributed to all BitLocker-protected devices using Group Policy Object (GPO) settings to ensure that only the administrator with a matching DRA certificate and private key can recover the information.

Administrators can use GPO settings to configure what recovery methods are required, disallowed, or made optional. For example, administrators can use GPOs to require that the recovery password for the OS drive is stored in AD. Administrators can also use GPO settings to determine whether a recovery password can be saved to a file on disk, printed, or viewed as text.

Select an Easy Deployment Method

In large IT environments, you can use a Microsoft script called EnableBitLocker.vbs to automate BitLocker deployment and configuration. This script calls on the Windows Management Instrumentation (WMI) providers for BitLocker and TPM administration. You can use the script as is, or you can customize it to better meet your organization's needs.

To run the script, you can leverage a startup script that is applied using GPO settings or a software distribution tool, such as Microsoft Systems Management Server (SMS) or System Center Configuration Manager (SCCM). The EnableBitLocker.vbs sample WMI deployment script can be downloaded from the MSDN BitLocker Deployment Sample Resources page at code.msdn.microsoft.com/bdedeploy/Release/ProjectReleases.aspx?ReleaseId=3205.

> BitLocker is a very powerful security technology that has reached a good level of maturity in Server 2008 R2 and Windows 7.

Prior to deploying BitLocker protection for an OS drive, you might need to check the disk partitioning on the target systems. On an OS drive, BitLocker requires a separate and active system partition. This is an unencrypted partition that contains the files needed to start the OS. In Windows 7, a separate active system partition is created automatically as part of the Windows installation process. On systems that were upgraded from a previous Windows version or on systems that come preconfigured with a single partition, the BitLocker setup wizard will automatically reconfigure the target drive for BitLocker by creating the separate and active system partition.

You don't want to use the BitLocker setup wizard for preparing hundreds or thousands of your systems with a single partition configuration for BitLocker. In those cases, you will want to use the Microsoft WMI script to enable BitLocker. Before you get there, you can also use a special tool called the BitLocker Drive Preparation Tool (BdeHdCfg.exe), provided by Microsoft, to prepare the systems' drives for BitLocker. You can find more information about this tool on the BdeHdCfg.exe Parameter Reference page at technet.microsoft.com/nl-be/library/ee732026%28en-us,WS.10%29.aspx.

For small BitLocker deployments, I advise you to use the BitLocker command-line tool Manage-bde.exe to configure BitLocker. This tool is designed to enable BitLocker on one computer at a time and to assist with administration after BitLocker is enabled. Again, before you use Manage-bde.exe to enable BitLocker on an OS drive, you might need to prepare the hard disk for BitLocker by running the BitLocker Drive Preparation Tool.

Get It Done the Right Way

BitLocker is a very powerful security technology that has reached a good level of maturity in Server 2008 R2 and Windows 7. It requires careful planning and a design that pays special attention to selecting the correct unlock method, defining a solid recovery strategy, and choosing an easy deployment method. With these three steps in mind, you're well on your way to BitLocker confidence.

InstantDoc ID 129258

Jan De Clercq
(jan.declercq@hp.com) is a member of HP's International Expertise Team and focuses on architecture for Microsoft-based IT infrastructures, identity management, cloud computing, and security. He's the co-author of *Microsoft Windows Security Fundamentals* (Digital Press).

VMware vSphere PowerCLI

Better management for your VMware servers

by Alex K. Angelopoulos

VMware vSphere PowerCLI extends PowerShell with cmdlets specifically for managing VMware servers. PowerCLI includes more cmdlets than PowerShell itself ships with, giving administrators broad access to VMware's internals. The downside of this wealth of tools is that it's easy to get lost inside PowerCLI if you don't start with specific goals in mind.

In this article, I walk through an entire PowerCLI working session that includes connecting to a server, performing some basic monitoring and management tasks, and disconnecting. The steps for using PowerCLI are always basically the same, so my example not only provides you with a framework for understanding PowerCLI but also demonstrates essential PowerCLI features. I also direct you to some of the most useful resources for discovering how to use PowerCLI for more specific tasks.

Prerequisites

To manage a VMware environment with PowerShell, your environment must meet three requirements. First, you need to have a VMware server environment available. VMware PowerCLI works with VMware ESX or ESXi 3.0 and later, as well as any VMware Server, VMware vCenter Server, or VMware vSphere product, version 2 or later.

Second, you need a workstation with PowerShell 2.0 installed; this is where you'll install PowerCLI. As of press time, the current version of VMware vSphere PowerCLI is 4.1.1. To get the installation package, go to the VMware vSphere PowerCLI homepage (www.vmware.com/go/powershell). Alternatively, you can go to the VMware Download Center (downloads.vmware.com) and search for PowerCLI. If you haven't already registered with VMware, you'll need to do so—the download is free. Please be aware that there's a similarly named but distinct product: vSphere CLI. Although vSphere CLI also allows command-line VMware management, vSphere CLI is a traditional console application, not the PowerShell toolkit. Make sure you download PowerCLI.

The third requirement is that you need access to the VMware server from the workstation you'll use for management. This might seem self-evident, but I mention it separately because there's minimal work for remote access to a VMware server. If you have access to the VMware server's network via VPN or similar technology, there's nothing specific you need to configure for access; you simply need to connect to the server's network. If you don't have VPN access but can configure port forwarding on the router for the VMware server's network, log on to the router and forward a port on the public side of the router to the VMware server's management port (443 by default). You can then use the public IP address or DNS name of the network and the public port number to connect to the server.

Loading PowerCLI

In general, you don't need to worry about the details of loading PowerCLI. The VMware PowerCLI installer provides you with a startup shortcut on the Start menu under VMware, VMware vSphere PowerCLI. When you start PowerShell from that shortcut, PowerShell automatically runs a configuration

Figure 1: Entering PowerShell credentials

script that loads the PowerCLI snap-in (among other things). If you always use the VMware PowerCLI shortcut to start a PowerCLI session, you can skip to the next section.

If you aren't using the shortcut, you can load the PowerCLI snap-in directly at a PowerShell prompt; you just need to know its name (i.e., VMware.VimAutomation .Core). Use this name with PowerShell's Add-PSSnapin cmdlet, like this:

```
Add-PSSnapin VMware.VimAutomation.Core
```

PowerShell will automatically find the snap-in and load it for you. This works from a script as well. If you try to use this command on a machine without PowerCLI installed, you'll get an error message similar to this: *The Windows PowerShell snap-in "VMware.Vim.Automation.Core" is not installed on this machine.* (As an aside for scripters, after running this command you can see if it succeeded by checking the value of the $? variable; it will be false if the snap-in didn't load.)

If you're loading the snap-in yourself instead of using the PowerCLI shortcut, I suggest that you at least glance at the configuration script. This script, Initialize-PowerCLIEnvironment.ps1, is in PowerCLI's installation folder. Beyond loading the snap-in and customizing the PowerShell interface, the configuration script also defines several PowerCLI-specific aliases and functions.

Connecting to a VMware Server

After you start the Power CLI session, use the Connect-VIServer cmdlet to connect to the VMware server you're managing. This step is typically very simple; you just use the Connect-VIServer cmdlet with the server name or address:

```
Connect-VIServer -Server
    192.168.1.21
```

PowerShell will then prompt you for credentials to let you connect to the VMware server using a standard PowerShell credentials dialog box such as the one that Figure 1 shows.

There's one situation in which this approach might not work. Suppose you're tunneling through the Internet to the remote VMware server. The best standard solution to this problem isn't to do something with PowerCLI but to set up a secure VPN to the remote site. If that's not an option, however, you can handle it with some router reconfiguration and the Connect-VIServer cmdlet.

Configure the remote site's router to forward an available port on its Internet side to port 443 on the VMware server on the private side. Connect-VIServer lets you specify an alternative port for situations like this.

As an example, suppose the VMware server is on a remote LAN. Because that LAN also has a mail server, there's a well-known name we can use to get to the router: mail.net.test. After connecting to the router and setting it to forward port 51234 on the outside to the VMware server's port 443, we'd use Connect-VIServer like this:

```
Connect-VIServer -Server mail.net.test
  -Port 51234
```

At this point, we're ready to actually do something on a VMware server. In the next few sections, we'll examine some basic tasks that are particularly useful or that provide a foundation for understanding more of the VMware infrastructure.

Checking Health and Performance

PowerCLI has a handful of cmdlets for quickly inspecting server logs and current statistics. The cmdlet for accessing logs is simply named Get-Log. However, you'll typically need to use the Get-LogType cmdlet first.

The reason for this is that VMware servers don't have a static set of logs. ESX and ESXi servers have a restricted set of logs because there are fewer core services on them compared with a full vSphere server, and even those logs exist in multiples because VMware rotates log files and retains backups. In the simplest case, you'll have a short collection (like the one in Figure 2, from an ESXi server). To get only the hostd log entries, you'd use the command

```
Get-Log -Key hostd
```

PowerCLI returns this as a blob of entry values, so to see the literal text in the hostd log, you need to expand the returned Entries value. You can expand Entries like this:

```
(Get-Log -Key hostd).Entries
```

In fact, if you really want to get all the existing log entries, no matter how many logs or entries there are, you can pipe Get-LogType into Get-Log and expand the entries for each log:

```
Get-LogType | Get-Log | ForEach-Object
  { $_.Entries}
```

However, this can be very resource-intensive, particularly for the server—and especially if you're looking at an ESX or ESXi server. If you won't be prefiltering the log entries and are really using this

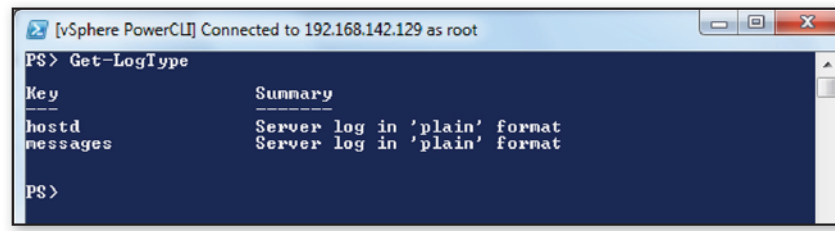

Figure 2: ESXi server logs

approach to get an idea of what's going on with a server overall, it might be better to use Get-Log's bundling feature. It will take as long as displaying the logs interactively, but you only have to do it once. Furthermore, this technique will also give you a complete set of server configuration files. All you need to do is run Get-Log with the Bundle parameter and specify a path on your local machine where you want to save the logs. This path must already exist; if you want to save the log bundle to C:\tmp\Server1, you must create the folder C:\tmp\Server1 yourself. The command you'd use looks like this:

```
Get-Log -Bundle -DestinationPath c:\
  tmp\Server1
```

When PowerCLI finishes bundling the logs and saving them, it returns a filename for the archive where they're stored. This is a standard zip file if you're using vCenter. Because ESX and ESXi servers use standard Linux infrastructure tools, the archive will be a tarred and gzipped file, with a .tgz extension; you can open these archives with WinZip or 7-Zip.

The log files give you information about significant events (and about machine configuration, if you use the Bundle feature). In contrast, you can use the Stat cmdlets to get statistics that give you a feel for system performance.

VMware servers automatically monitor a range of performance statistics for servers, virtual machines (VMs), resource pools, clusters, and hosts. We'll look at the non-server entities in the next section because you need to get references to those entities to look at their statistics. For the currently connected server, however, you don't need to do anything complicated to find out what statistics VMware tracks. To see the statistic types, just enter the command

```
Get-StatType
```

This command returns a list of the names or MetricIds of statistical quantities that VMware tracks. This list can be overwhelming, particularly because the list can contain duplicates. You can get a clean listing by using

```
Get-StatType | Sort-Object -Unique
```

Figure 3 shows the initial output from this command; although still lengthy, the list is abbreviated and alphabetized. You can use one or more of these names, or a wildcard variant with Get-Stat, to retrieve server statistics. Either of the following is a legitimate statistic name specification:

```
Get-Stat -Stat cpu.usage.average
Get-Stat -Stat cpu.u*.average
```

Get-Stat is useful without enumerating statistic types, however. You can simply enter

```
Get-Stat
```

to retrieve core CPU, disk, memory, and network-usage statistics for the server.

Listing Virtual Machines

You can use a simple command to register VMs on a VMware server. To retrieve every VM registered on the currently connected server, whether running or not, use the command

```
Get-VM
```

This command returns a list of VMware machine objects; the default display will show each machine's name, PowerState, CPU count, and memory size. The Get-VM cmdlet supports several parameters you can use to filter the output. The one I most commonly use is the Name parameter. You can specify single or multiple names, possibly containing wildcard values. The following command will return all the VMs whose names begin with either XP or Vista:

```
Get-VM -Name XP*,Vista*
```

The Name parameter is useful for returning all machines in an installation, or a handful based on name pattern. However, if you work with a sizable VMware infrastructure, you should explore the other parameters for Get-VM in the Help documentation; you can filter the data returned against one or more specific data stores, data centers, folders, or clusters.

Creating New Machines, Clones, and Snapshots

Machine population management is a huge topic. However, some basic cases are simple to demonstrate (and are in the Help documentation as well, along with cases that are more complex).

When setting up new production VMs, I usually start from a template. All you need is the template name and name for the new machine. Working from the Windows XP Professional VMware template named XppTemplate and creating a new guest named Xpp05 works like this:

```
New-VM -Template XppTemplate -Name
  Xpp05
```

Cloning uses the New-VM cmdlet as well, only you specify an existing VM rather than a template. This will be a full copy of the original system (there's no way to create a linked clone via cmdlets). To clone the VM named Xpp05 to a new test system named XpTesting, use this command:

```
New-VM -VM Xpp05 -Name XpTesting
```

Finally, to take an immediate snapshot of XpTesting so that I can roll back to it if necessary, I use New-Snapshot:

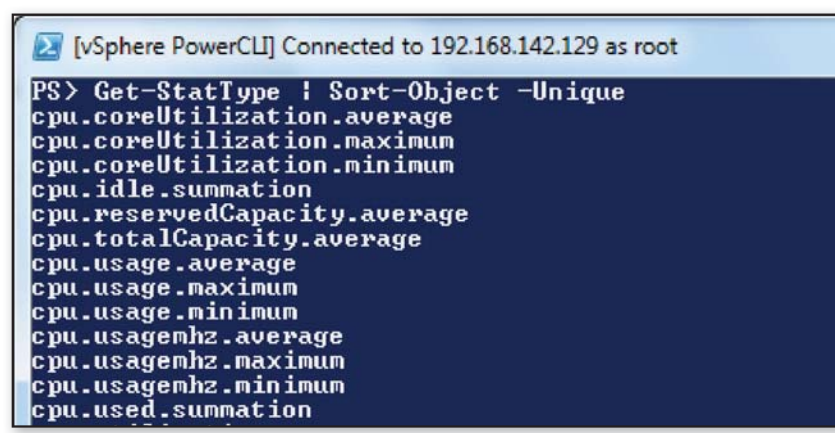

Figure 3: List of statistic types

```
New-Snapshot -VM XpTesting -Name
  Creation
```

This command creates a snapshot named Creation.

Cycling Virtual Machines

VMware has cmdlets for managing machine states. For most of them, you can only perform the tasks gracefully if you installed VMware Tools within the VM.

To start a VM, use the Start-VM cmdlet. You can start the machine Xpp05 like this:

```
Start-VM -VM Xpp05
```

The other operations are where VMware Tools makes a difference. Without VMware Tools installed on the guest OS, you would restart, suspend, or shut down Xpp05 with the following commands:

```
Restart-VM -VM Xpp05
Suspend-VM -VM Xpp05
Stop-VM -VM Xpp05
```

If you do have VMware Tools installed on the guest, VMware can tell the OS to initiate the changes instead of forcing them. To restart, suspend, or shut down the commands via VMware Tools, you would use the following commands:

```
Restart-VMGuest -VM Xpp05
Suspend-VMGuest -VM Xpp05
Shutdown-VMGuest -VM Xpp05
```

Be warned: Although the VMGuest cmdlets are polite to the guest OS, they don't warn connected users about state changes. If you shut down running machines with any of these cmdlets, an interactive user will either see what looks like an unstoppable OS shutdown (with Shutdown-VMGuest) or will simply see the session disappear (if using Stop-VM).

Moving Running Virtual Machines

If you have multiple VMware servers, it's possible to move running VMs in correctly configured shared storage from one server to another using VMware VMotion. Suppose you have a VM named XP17 on VMware host VH1. You have a connection to server VH1, and you need to move XP17 to VMware host VH2. All you need to do is get the VM object, then use that object and the target host address to make the move, like this:

```
Move-VM -Destination VH2 -VM XP17
```

Suppressing Confirmations for PowerCLI Cmdlets

If you've stepped through at least a few of the PowerCLI examples here, you've probably noticed that when you do something that could affect either your connection to a VMware server or a user's session on a machine, the PowerCLI cmdlet you're using will give you a confirmation prompt. This can get irritating in scripts if you need them to run even partially unattended.

For any cmdlet that has confirmation prompts, you can suppress the prompts if you explicitly set the Confirm parameter to false. You need to know one detail in order to use Confirm this way; you must join the parameter name to the $false value with a colon (:) so that it looks like -Confirm:$false. For example, to suppress Restart-VMGuest's confirmation prompt, you need to use it like this:

```
Restart-VMGuest -VM Xpp05
  -Confirm:$false
```

This is necessary because the Confirm parameter is a switch parameter. The name by itself with no argument tells PowerShell that you want a confirmation prompt. By using the colon you explicitly tell PowerShell that the value after the colon is an argument to the parameter.

Disconnecting from the Server

A last step that's valuable in scripts—and in interactive PowerCLI sessions if you're trying to work as cleanly as possible—is to shut down the server connection explicitly. Just use

```
Disconnect-VIServer
```

When you do this, the current server connection will be closed.

Although PowerShell and VMware will eventually clean up connections anyway without this step, I strongly recommend getting in the habit of using it, particularly in scripts designed to perform complete tasks on a server. Explicitly closing the connection does release resources immediately and reduces the risk of accidentally misapplying operations to the wrong server.

Getting Help: PowerCLI Resources

Although we've stepped through a complete VMware administration session, we've touched on only a handful of the simpler cmdlets you can use and very few of the tasks you can accomplish in PowerCLI. Several resources can help you get up to speed with using PowerCLI quickly.

PowerCLI's cmdlets are documented internally just as the native PowerShell cmdlets are, so you can use Get-Help and Get-Command to explore the cmdlets in general or specifically from within a PowerCLI session. PowerCLI also provides the analogous PowerCLI-specific commands Get-VICommand and Get-PowerCLIHelp.

The standard PowerCLI installation includes references that you can use outside of PowerCLI. If you look in the Start menu's VMware vSphere PowerCLI folder, you'll find a link to a Windows Help file called *vSphere PowerCLI Cmdlets Reference*; this is a searchable graphical interface version of the console-based cmdlet help. In the same location you can also find the PDF-based *vSphere PowerCLI Administration Guide*, which provides an extended walkthrough of using PowerCLI.

Finally, there's a dynamic and extremely helpful PowerCLI user community. You can find the community online at communities.vmware.com/community/vmtn/vsphere/automationtools/powercli. PowerCLI itself also contains a Get-PowerCLICommunity cmdlet that automatically opens a browser window on this community site. Both user-community experts and VMware developers participate in the community's web-based discussions; the site also features a growing collection of scripts and other useful PowerCLI documents.

InstantDoc ID 129623

Alex K. Angelopoulos (aka@mvps.org) is an IT consultant, an MCSE, and a contributing editor for *Windows IT Pro*. As an avid scripter, he regularly writes about administrative automation using WSH, PowerShell, and related technologies.

Virtual Desktop Infrastructure, Part 2: Finally, VDI

Create a VDI solution based on Windows Server 2008 R2

by John Savill

In "Virtual Desktop Infrastructure, Part 1" (January 2011, InstantDoc ID 129007), I covered the technologies that make up desktop virtualization, including application virtualization, roaming profiles, folder redirection, and OS virtualization. I explained how combining these technologies lets you dynamically construct the entire user environment as needed on a local OS, a presentation virtualization platform such as Remote Desktop Services (RDS), and a Virtual Desktop Infrastructure (VDI) environment. In this article I discuss the components necessary to create a pure Microsoft VDI solution based on Windows Server 2008 R2.

Prerequisites

Before we step through the typical process of connecting to a hosted desktop, it's helpful to identify the types of services necessary for a successful VDI architecture. Figure 1 shows the major steps required for VDI functionality, from the initial user contact all the way to a usable VDI session—except for technologies such as Microsoft Application Virtualization (App-V), roaming profiles, and so on, that actually populate the OS environment.

The process involves seven distinct steps. First, users must find the remote desktops they can connect to, which can include presentation virtualization sessions (Terminal Services), published applications, and VDI sessions. Although an RDP file can be created and deployed to users through various methods, a more dynamic approach is to use the Remote Desktop Web Access role service, which presents a browser-based list of available connections from which the user can choose. An alternative to users having to browse a website and have a list of dynamically published connections and applications is to use the Windows 7 *RemoteApp and Desktop Connections* feature to subscribe to an RSS-based feed from the Remote Desktop Web Access servers to automatically populate the available connections, which can be configured using a Group Policy–deployed script. Users can then see these published services straight from their Start menu.

The second step is to create a list of published applications and connections that are presented to the user. To accomplish this task, the Remote Desktop Web Access server communicates with the Remote Desktop Connection Broker, which has knowledge of the VDI pools, personal desktops, and other published connections and applications through its own communication with configured RemoteApp sources.

In the third step, the Remote Desktop Connection Broker communicates with Active Directory (AD) to determine the user's exact access. This communication also exposes any personal desktop configurations, which I discuss later in the article.

No matter what method is used, whether Remote Desktop Web Access, the *RemoteApp and Desktop Connections* feature, or a deployed RDP file, users now have an RDP file that can be used to initiate the connection (see step 4). If a user were outside the corporate network, then a direct RDP connection would be blocked by most organizations' firewalls. Thus, we'd need to initiate a secure VPN connection. However, we have an alternative solution that doesn't require any end-user action

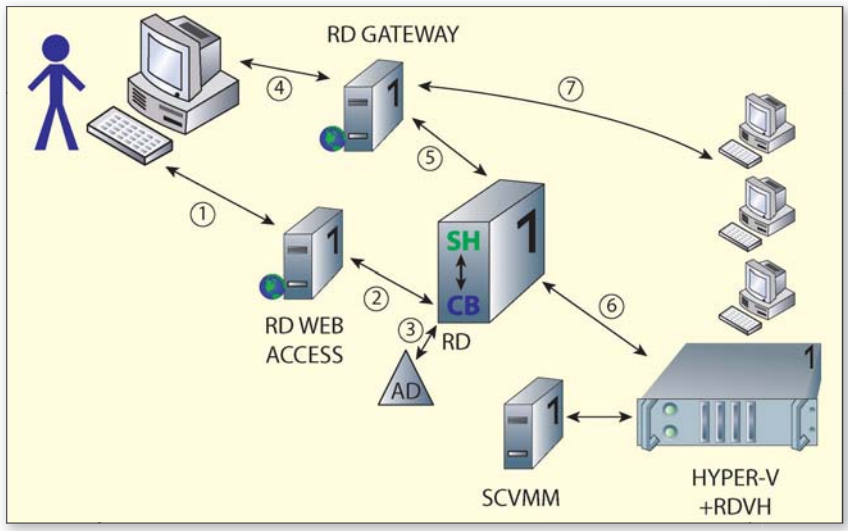

Figure 1: VDI components

or additional client-side software. Windows Server 2008 introduced Terminal Services Gateway, which allows the RDP traffic to be encapsulated in HTTPS (port 443) packets. (Terminal Services Gateway is renamed Remote Desktop Gateway in Server 2008 R2.) With Remote Desktop Gateway in the architecture, we place the Remote Desktop Gateway server in the demilitarized zone (DMZ), or more securely behind some kind of firewall or proxy. Clients then connect to the RDP destination through the Remote Desktop Gateway by adding the Remote Desktop Gateway server as part of the RDP file configuration that's given to the client. The client encapsulates the RDP traffic in HTTPS and sends the HTTPS-encapsulated RDP data to the Remote Desktop Gateway, which extracts the RDP data and forwards it to the RDP destination. When traffic comes back from the RDP destination bound for the client, the Remote Desktop Gateway encapsulates the RDP data in HTTPS and sends the HTTPS-encapsulated RDP data to the client. With this technology, users outside the corporate network can still access all RDP resources without additional steps or software. Users on the corporate network would bypass the Remote Desktop Gateway and communicate directly with the RDP destination. Authentication to the Remote Desktop Gateway can be through traditional Windows authentication or via smart card, including the ability to connect through current logon credentials, which prevents having to reenter credential information.

In step 5, the user needs an initial RDP connection point because the VDI client virtual machine (VM) destination isn't yet known, unless the user has a personal desktop configured. A Remote Desktop Session Host is configured in redirection mode, which means the server acts as the connection point for the user's RDP connection, then redirects the client to the true endpoint, the VDI session. The Remote Desktop Session Host communicates with the Remote Desktop Connection Broker to determine what the RDP target should be for the requesting client. The Remote Desktop Session Host and Remote Desktop Connection Broker are on the same host in this architecture, but you don't have to place them on the same host (although it's common to do so because of how closely the two roles work together).

In step 6, the Remote Desktop Connection Broker communicates with the Remote Desktop Virtualization Host role service that's enabled on the Hyper-V boxes to check the state of the VMs, start the VM if required, and gather any needed information, such as IP address of the client VM OS, which is then passed back to the Remote Desktop Connection Broker, to the Remote Desktop Session Host in redirection mode, then back to the client.

In step 7, the client makes an RDP connection to the destination client VM via the Remote Desktop Gateway (if connecting from outside the corporate network); the connection is now complete. At logon, the profile will be downloaded through roaming profiles, along with data available via folder redirection and applications available through App-V.

Although the process of connecting to a hosted desktop includes a lot of components and steps, it's a smooth and near-instant process from the users' perspective, providing a full-featured client experience from any RDP client.

One item that Figure 1 includes but that I didn't discuss is Microsoft System Center Virtual Machine Manager (VMM). Although VMM isn't 100 percent necessary for a VDI environment, it can be helpful in managing your VDI environment by providing Hyper-V farm management capabilities, libraries, PowerShell support, and much more.

For a pure VDI environment, you can use the free Microsoft Hyper-V Server solution and save the cost of purchasing Server 2008 R2. Server 2008 R2 SKUs provide the same Hyper-V capabilities as the free Hyper-V Server solution but have additional capabilities beyond virtualization. In addition, Server 2008 R2's Hyper-V has additional server virtual instance rights—which we don't need, so you should save your money and use the free solution instead.

When we look at all the steps involved in the connection process, we essentially have five RDS roles. If you decide to use RemoteFX with SP1 (which I'll cover in a future article), that makes six Remote Desktop role services. Using any of the Remote Desktop role services requires clients to have RDS CALs, in addition to any virtual desktop access/software assurance rights. Let's dive a little deeper into each of the services and any special considerations related to VDI. For full VDI deployment information, refer to the Microsoft TechNet page "Getting Started: Remote Desktop Services" at technet.microsoft.com/en-us/library/dd736539(WS.10).aspx.

Remote Desktop Web Access

Remote Desktop Web Access provides the initial entry point, giving users a web-based interface to select the desired VDI or published desktop/application target. Although not absolutely required, Remote Desktop Web Access helps give a simple-to-use portal that supports forms-based authentication and single sign-on (SSO), in addition to differentiating between public and private computers.

Figure 2: Web-based and *RemoteApp and Desktop Connections* views of published applications and desktops

Windows 7 introduces the *RemoteApp and Desktop Connections* feature, which—although not directly part of Remote Desktop Web Access—allows a feed from Remote Desktop Web Access to populate the same content shown in the website directly into the Start menu, thus avoiding the need to use the website. Figure 2 shows both the local Start menu and the website view. Note that we can see not only the VDI pool but also published applications from remote desktop session hosts.

If you want to create a configuration file that automatically configures the *RemoteApp and Desktop Connections* feature for Windows 7 clients, use the Create Configuration File option in the Remote Desktop Connection Manager. This option creates a file that you can run on a Windows 7 client to automatically configure the client. It's a great solution for large deployments.

Remote Desktop Connection Broker

Server 2008 R2's Remote Desktop Connection Broker role service is one of the major components that allows an all-Microsoft VDI solution, giving RDS the ability to balance and track connections to non–terminal servers—specifically, the ability to manage connections to client OSs. In addition, Server 2008 R2 introduces the ability for the Remote Desktop Connection Broker role service to balance remote applications and support servers with different published applications, allowing a sum view of all the different applications gathered from all servers in the farm, to be displayed to the user—thus removing the need for all servers to have exactly the same applications.

The Remote Desktop Connection Broker role service is really the brains of the VDI environment. It communicates with and controls the other components, working particularly closely with the Remote Desktop Session Host in redirection mode, which is why the Remote Desktop Connection Broker and Remote Desktop Session Host in redirection mode are frequently placed on the same OS instance. However, when you start having more than 250 simultaneous connections, you might need to consider breaking the roles onto separate servers.

VDI pools are created through the Remote Desktop Connection Broker role service—specifically through the Remote Desktop Connection Manager administrative tool, which manages the published application, published session, and virtual desktop resources. The Remote Desktop Connection Manager is a tool you'll come to love when managing your VDI environment. This single tool not only lets you create VDI pools but also configures the Remote Desktop Virtualization Host, Remote Desktop Web Access, and Remote Desktop Session Host role services. The tool also places the specified Remote Desktop Session Host in redirection mode, saving you the overhead of manually configuring each service for its VDI role.

When we create a VDI pool, we specify the Hyper-V servers that are hosting the VMs, then select the client VMs that will be part of the pool. The VDI pool creation wizard takes care of most of the other necessary configuration.

The Remote Desktop Connection Broker role service handles all incoming VDI requests and first checks to see if the user has a disconnected session within the VDI pool. If so, the user is reconnected; otherwise, a currently unused virtual desktop is assigned to the user.

Remote Desktop Session Host in Redirection Mode

The concept of using a Remote Desktop Session Host in redirection mode isn't new. We've used it with previous version of Windows Server when we had a large terminal server farm. To prevent users from having to connect to different terminal servers, the initial entry point is always a designated session host that does nothing more than talk to the broker and redirect the RDP connection to the correct RDP endpoint. The same process occurs in a VDI environment; we still need an initial RDP connection point for the RDP clients, which is exactly what the Remote Desktop Session Host in redirection mode provides. It then redirects the RDP client to the correct client OS VM that will provide the desktop OS.

The Remote Desktop Session Host is what we used to think of as a terminal server: a server that provides the hosting of sessions. But because Server 2008 R2 has so many RDS parts, we now need to be a little clearer when we actually have a server providing sessions—hence the term Remote Desktop Session Host.

We can manually configure a Remote Desktop Session Host to be in redirection mode. However, if we use the Remote Desktop Connection Manager VDI pool creation wizard, the Remote Desktop Session Host configuration is done automatically, as Figure 3 shows. Note that once a Remote Desktop Session Host server is placed in redirection mode, it can't also be used to host regular sessions; it can only perform redirections.

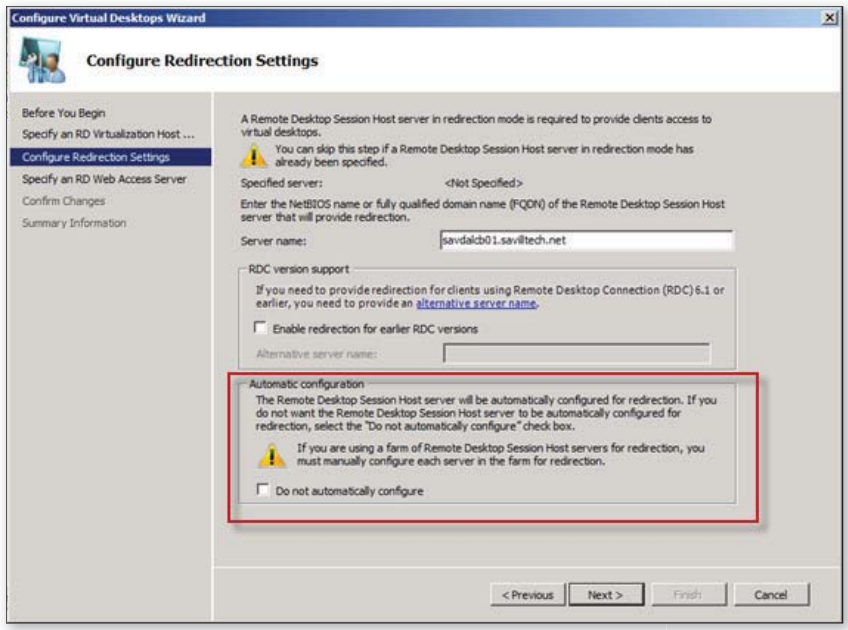

Figure 3: Automatic Remote Desktop Session Host configuration through the VDI pool wizard

Remote Desktop Virtualization Host

The Remote Desktop Virtualization Host role service is installed on any Hyper-V host that will be participating in a VDI pool. This role service lets the Remote Desktop Connection Broker role service communicate with the Hyper-V hosts, start and stop VMs, and gather internal information to enable client connections.

Remote Desktop Gateway

The Remote Desktop Gateway role service allows RDP traffic to be encapsulated in HTTPS packets, enabling secure RDP connection through corporate firewalls without having to open up firewall ports or use additional VPN solutions. We can use Remote Desktop Gateway to configure who can connect through the Remote Desktop Gateway service, what they can connect to, and the supported RDP settings, such as device redirection.

Highly Available VDI

Implementing VDI involves centralizing the desktop environment into the data center. If a server fails, users can still use their local OS environment. But when we implement VDI, the entire desktop is in the data center—so if part of the VDI environment fails, the user loses the entire desktop environment and is unable to work. Ensuring that all the elements of a VDI environment are fault tolerant is crucial. Fortunately, we have a solution for each part of the VDI architecture.

Looking at the Remote Desktop Connection Broker role service as the brains of the VDI environment, we use failover clustering because the role service is cluster aware, which ensures that it's always available and survives the loss of any node. Because we normally co-locate the Remote Desktop Session Host in redirection mode with the Remote Desktop Connection Broker role service, you need to ensure that you install the Remote Desktop Session Host and configure it in redirection mode on all nodes in the cluster. We create a DNS Resource Record for the session hosts, which has the IP address of each session host configured. Thus, when clients connect, all the IP addresses are sent to them in varying order. If one server isn't available, then the clients will just try the next IP address.

The Network Load Balancing (NLB) service is used to balance between multiple Remote Desktop Gateway instances. We can use the NLB service that's part of Windows or we can use a hardware load balancer. We can use the same NLB technologies to make Remote Desktop Web Access highly available.

Hyper-V hosts can be placed into failover clusters so that if a Hyper-V host fails, client VMs can be moved to other hosts. Live Migration can be used to move running client VMs between hosts for maintenance purposes if necessary. However, because the user state and data are separate from the OS instance, if an OS instance is lost, users just reconnect to an alternative OS. User states and data are then available again.

Personal vs. Pooled Desktops

So far I've discussed pools for VDI clients, which is a configuration in which several VMs running the client OS are grouped together into a pool. As users connect to the pool, they're automatically assigned one of the VMs that isn't currently in use. After the user logs off, the VM is placed back into the pool. Because a user potentially (and probably) gets a different VM each time, we need to have various desktop virtualization solutions in place (e.g., roaming profiles, folder redirection, application virtualization) to ensure a consistent desktop experience.

Pooled desktops should be the default for all users. However, certain users might need the same client OS instance every time they connect. Maybe they're modifying the OS in some way, or perhaps they have an application that needs to be installed because it can't be virtualized. Whatever the reason, we have the capability to statically assign a VM to a particular user so the user always gets the same client OS. This is known as a personal desktop and is configured through the Microsoft Management Console (MMC) Active Directory Users and Computers snap-in, as Figure 4 shows. A user can be assigned only one personal desktop, a VM can be assigned to only one user as a personal desktop, and a personal desktop must not be in a VDI pool but should just be a regular VM in the environment. Make sure the personal desktop name exactly matches the name of the VM, which needs to be the Fully Qualified Domain Name (FQDN)—for example, client.savilltech.net—which means you need to name the VMs the FQDN of the client OS instance.

Client VM Configuration

We'll want to use Windows 7 as the client OS within our VDI VMs for the best experience. With an RDS-based VDI deployment, we must create all the VMs in advance, install the OS, and add to the pool—there's

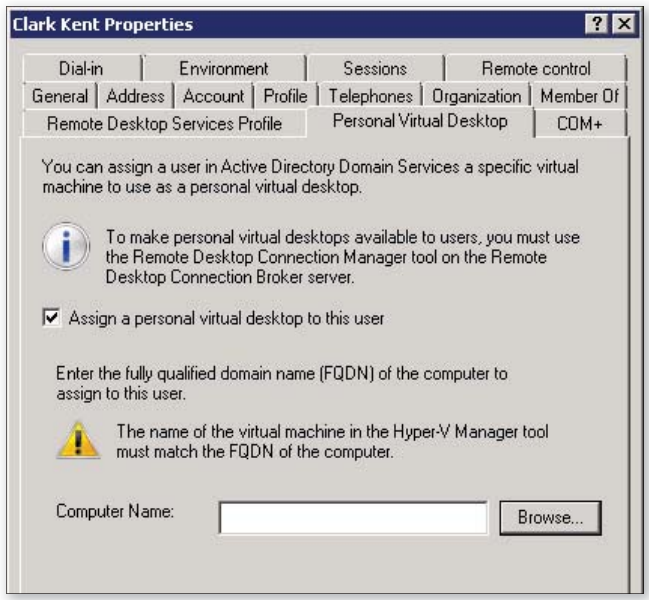

Figure 4: Assigning a personal desktop to a user

no dynamic creation or streaming of the OS to populate the VM. (This capability comes with Citrix XenDesktop, which I'll cover in a future article.) We can use technologies such as VMM to simplify the VM and OS deployment process.

We need to perform some particular configuration inside the client VMs to enable management of the client OS by the Remote Desktop Virtualization Host and connectivity from the clients. The major steps you need to perform in each client OS are:

- Enable Remote Desktop
- Add connecting users to the Remote Desktop Users group
- Allow RPC
- Configure firewall exceptions for managing remote services
- Configure RDP permissions for the Remote Desktop Virtualization Host

You can manually perform all the required steps for each VM (see Microsoft TechNet's *Remote Desktop Services in Windows Server 2008 R2,* Appendix A, "Configuring the Virtual Machine Manually," at technet.microsoft.com/en-us/library/ff603843(WS.10).aspx); you can use Group Policy to automate the steps to apply the configuration; or you can use a Microsoft-provided script to simplify the entire process (available from gallery.technet.microsoft.com/ScriptCenter/en-us/bd2e02d0-efe7-4f89-84e5-7ad70f9a7bf0).

There's another interesting aspect to the client VM OS environments. We have multiple client OS environments that are shared by multiple users. We should be locking down these shared client environments so that users can't change the OS configuration and add or remove software. However, depending on the environment, we might want the OS to be reset to a known state each time a user logs off. We can use Hyper-V and RDS's snapshot capabilities to achieve this goal.

For each VM, we can create a snapshot that includes RDV_Rollback in the snapshot name. This snapshot should be taken when the VM is in a clean state because you want the OS to be reset each time a user logs off. The snapshot can be taken when the VM is running or not running, but you must make sure no one is logged on when you take the snapshot. When a user logs off from a VDI VM that was connected to via the Remote Desktop Connection Broker role service, the VM is reset back to the RDV_Rollback snapshot. Note that this RDV_Rollback capability applies only to VMs in a pool, not to personal desktops.

If you choose to use RDV_Rollback to ensure that each user gets a clean OS environment, a wrinkle is introduced into the VDI environment related to AD domain membership. Typically, the computer account password of the OS instance automatically changes every 30 days. If we restore to a checkpoint periodically—for example, after each logoff—the old machine account password that's present in the RDV_Rollback will no longer be valid after the computer changes its password—which could happen while a user is logged on past day 30, which would cause subsequent logons to fail and require an account reset. There are several options to solve this problem. One solution is to disable the machine account password change, which we can accomplish through Group Policy. However, this approach can have security ramifications and therefore isn't recommended.

Another option is to actually delete all the client VMs, recreate them, and create a new RDV_Rollback periodically (i.e., an interval less than the AD machine account password change interval). This approach might seem ridiculous, but you can use VMM and Microsoft's scripts for creating VMs for VDI (available at gallery.technet.microsoft.com/scriptcenter/en-us/904bd2c8-099d-4f27-83da-95f5536233bc). These scripts let you automate the bulk creation of VMs for your VDI environment, which makes it far easier to recreate your VDI VMs as needed and much more realistic to periodically recreate the VMs. Using this approach also solves another issue. If we keep the VMs around for a prolonged period of time, we need to patch them and perform regular desktop maintenance. If we recreate the VMs every 4 weeks, then all we need to patch is a single master image that's replicated to create all the VMs in the VDI pool.

Next Steps

RDS-based VDI is a great solution that you can realistically use in production environments. For some large deployments, the static nature of the client VMs might be prohibitive—which Microsoft's partnership with Citrix addresses and which I'll cover in the next article in this series. In the meantime, refer to the Microsoft TechNet page "Getting Started: Remote Desktop Services," at technet.microsoft.com/en-us/library/dd736539(WS.10).aspx, for step-by-step VDI instructions. Within a day, you could easily set up a VDI environment in your lab environment and start experimenting.

InstantDoc ID 129572

John Savill
(john@savilltech.com) is a Windows technical specialist, an 11-time MVP, and an MCITP: Enterprise Administrator for Windows Server 2008. He's a contributing editor for *Windows IT Pro*, and his latest book is *The Complete Guide to Windows Server 2008* (Addison-Wesley).

FEATURE

Use the MAP Toolkit for a Smooth Windows 7 Migration

Inventorying your network applications and drivers

by Greg Shields

Editor's Note: This article is an excerpt from Greg Shields's new book, *Automating Windows 7 Installation for Desktop and VDI Environments* (Realtime Publishers); reprinted with permission. You can download a free copy of the entire book at nexus.realtimepublishers.com/awidv.php?ref=winitpro.

Deploying Windows 7 is an involved process. To ensure a smooth migration, you should first inventory your network applications and drivers. This process begins with installation of the Microsoft Assessment and Planning (MAP) Toolkit. With it, I'll show you how to gather a report of the software that's installed on computers around your network. With that information in hand, I'll point you to Microsoft's Windows 7 Compatibility Center. This site is an online clearinghouse of applications and their compatibility status. You can compare the software in your report to that in the Compatibility Center to see which will work and which won't.

But that's not all the MAP Toolkit is good for. Drivers can be automatically injected into images as they're deployed, and it's useful to know exactly which drivers your computers will need. The MAP Toolkit can also collect that information for you if you know where to look.

Installing the MAP Toolkit and Collecting Inventory

Begin by downloading the MAP Toolkit from Microsoft's website (technet.microsoft.com/en-us/solutionaccelerators/dd627342) and installing it to the Windows Deployment Services (WDS) server. Using the MAP Toolkit requires first installing a copy of Microsoft Office 2007 SP2, as well as the .NET Framework. The MAP Toolkit will automatically install a copy of SQL Server Express to the computer as it begins its installation. After the MAP Toolkit is installed, you'll be asked to create a new inventory database.

Figure 1: MAP Toolkit console

Figure 1 shows what the MAP Toolkit console will look like after installation. You should immediately notice that the MAP Toolkit has far more capabilities than simply searching your network for installed software. Other assessments are available that help determine Windows Server roles that have been installed on servers, where SQL Server components have been deployed, and even where virtual machines (VMs) might be hiding on your network.

Inventorying the software in your environment starts by clicking the *Inventory and Assessment Wizard* link (which you can see in Figure 1). Clicking this link opens a wizard that you'll use to configure the types of inventory to be collected. Windows, Linux, VMware, Exchange Server, and SQL Server computers are

Learning Path

See these articles for more information about Windows 7 migrations:

"Windows 7 Under the Hood," InstantDoc ID 103218

"XP to Windows 7 Migration with Microsoft Deployment Toolkit 2010," InstantDoc ID 103607

"Application Migration with MDT 2010," InstantDoc ID 125462

"Perform Windows 7 Bare-Metal Installations with MDT 2010," InstantDoc ID 125154

"Create Windows 7 Media for Deployment," InstantDoc ID 104644

"Q. What's Microsoft P2V Migration for Software Assurance?" InstantDoc ID 128797

all options for inventorying. I'll be using only the Windows-based computers scenario, because this scenario provides the information I'll need for a Windows 7 upgrade.

The wizard's second page shows the multitude of methods the MAP Toolkit will use in discovering computers to inventory, as Figure 2 shows. Because my computers are all members of an Active Directory (AD) domain, I can select the first and second check boxes to find them. Other computers not on the domain can be discovered either via IP ranges, by entering in computer names manually, or via a text file.

Subsequent wizard pages provide locations to enter AD credentials, to restrict inventory to specific organizational units (OUs), and to add additional domains or workgroups if they are discovered by the tool. The page titled All Computers Credentials allows you to enter a list of possible credentials the tool can use in attempting to inventory discovered computers.

It is within the All Computers Credentials and Credentials Order pages where the MAP Toolkit truly shines. You can see in Figure 3 that I have entered credentials for two different domains: COMPANY and SPECIALIZED. Additional workgroups or specific computer credentials can be added as well. Doing so will give the inventory process plenty of username and password options as it authenticates to discovered computers.

Click Finish to complete the wizard and begin the discovery and inventory process. Be aware that this process can take a considerable quantity of time, particularly if your scope is large. Version 5.0 of the MAP Toolkit, the version used in this example, is reported to discover and inventory up to 100,000 computers. Gathering information from that quantity of computers, as you can imagine, is going to take a while.

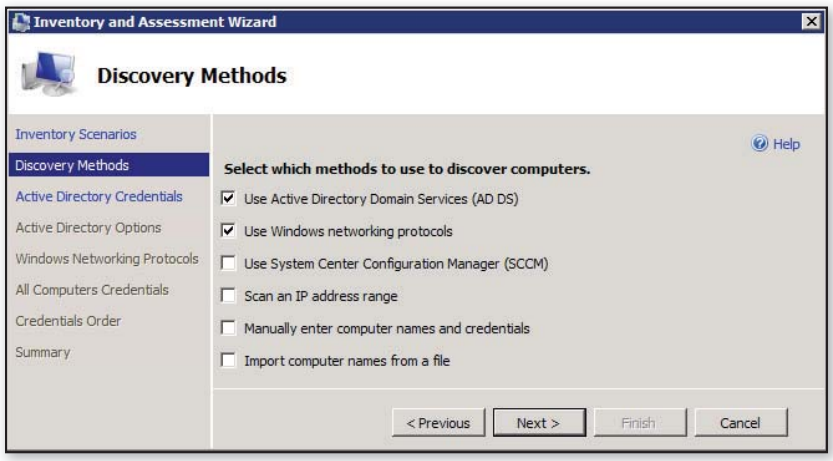

Figure 2: Selecting methods to discover computers

Figure 3: Setting an order for credentials

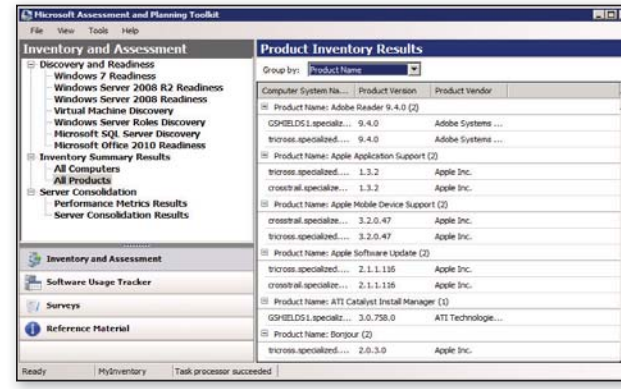

Figure 4: Product inventory results

Note that the MAP Toolkit's inventory process uses Windows Management Instrumentation (WMI) queries to gather its information. Ensure that the Remote Administration firewall exception has been enabled on any computers that will be queried by the MAP Toolkit.

Figure 4 shows a report of the products the MAP Toolkit found on my network. You can see that Adobe Reader 9.4.0 was discovered on two computers. A set of three Apple applications was found on another two, as well as an entire list of software from all sorts of vendors. This screen inside the report is relatively static, giving you little more than a view of the software that the MAP Toolkit has found inside your network.

A much more useful representation of the data found by the MAP Toolkit can be created by clicking the Windows 7 Readiness link in the *Inventory and Assessment* pane. The resulting Windows 7 Readiness summary provides some high-level information about the computers found in the discovery and inventory process. You can learn in this screen how many computers

Device Compatibility Summary

Source of Device Driver	Hardware Devices	Percentage
Included on the Windows 7 DVD	131	68 %
Available from Microsoft Update	1	1 %
Information available from manufacturer website	1	1 %
Contact the Device Manufacturer (unknown driver or incompatible)	61	31 %
Total	194	100 %

Figure 5: Device compatibility summary

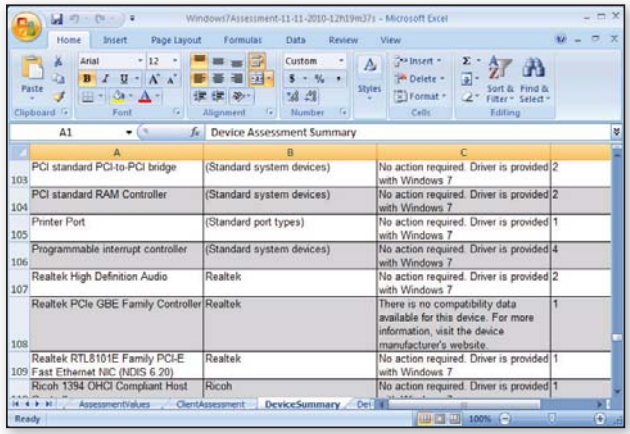

Figure 6: MAP Toolkit report's Excel spreadsheet

have hardware that is powerful enough to support Windows 7. You can also learn how many drivers your computers will need that are and are not included on the Windows 7 CD-ROM. Figure 5 shows a snippet of the summary screen. This screen tells me I'll need to locate manufacturer drivers for 61 of the 194 drivers my computers say they need.

Creating and Using MAP Toolkit Reports

In the right pane of the Windows 7 Readiness Summary Results page, click the link labeled Generate Report/Proposal to generate a report. Click View, then click *Saved Reports and Proposals* to open an Explorer window. In this window, you'll find a Microsoft Word document that contains some useful project planning information about your Windows 7 readiness.

You'll find even more useful information in the accompanying Excel spreadsheet. Inside that spreadsheet is detailed information about each inventoried computer, its hardware configuration, and any installed software and drivers. Figure 6 shows one of the tabs in that spreadsheet. In it you can see that at least one computer on my network reports it will need the Realtek High Definition Audio driver. Happily, that driver is available on the Windows 7 media, so I don't need to worry about it. Another computer reports it needs the Realtek PCIe GBE Family Controller, which isn't on the Windows 7 media. I'll need to locate that driver from its manufacturer's website and add it to my Out-of-Box Drivers node in my Microsoft Deployment Toolkit (MDT) deployment share.

By reviewing the drivers in this Excel spreadsheet, I now know which drivers I'll need to make available in MDT so that my images will deploy correctly. This report all by itself gives me the data I need to ensure that my deployment goes as smoothly as possible.

A second tab on this Excel spreadsheet gives me a punch list for tracking down the compatibility status of applications that are installed on my computers. That tab, labeled Discovered Applications, lists each application, its version number, and the number of instances found on the network during the last inventory pass.

I mentioned at the beginning of the article that Microsoft has created an online clearinghouse of application compatibility status information. That clearinghouse is called the Windows 7 Compatibility Center. Go to www.microsoft.com/windows/compatibility/windows-7/en-us/default.aspx to check out its constantly updated list.

I went to the website and ran a search on Adobe Reader. I already know from my MAP Toolkit report that I have two copies of Adobe Acrobat 9.4.0 on my network. As Figure 7 shows, running the search told me that Adobe Reader version 9 is compatible with Windows 7. It also told me that version 8 compatibility requires an action, specifically a free upgrade, which is useful information.

Combining this website with the information in my MAP Toolkit report, I can quickly identify which applications will work and which won't. For some, I might learn that they'll require a patch or some other special configuration to function.

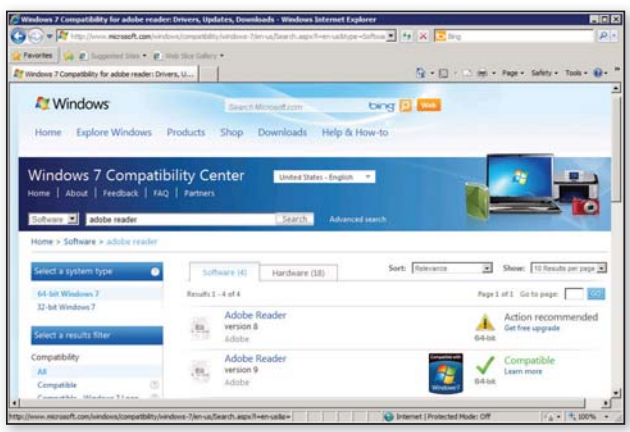

Figure 7: Using the Windows 7 Compatibility Center

Successful Migration

There really is more to deploying Windows 7 than just deploying Windows 7. Being successful with any migration or upgrade project entails knowing the drivers and applications that are on your network in comparison with those that won't work well atop the new OS. Microsoft's tools get you part of the way there. Your next, and arguably larger, task requires using this information to ensure software compatibility. Thankfully, with these tools in hand, answering your business's compatibility question won't be an impossible job.

InstantDoc ID 129578

Greg Shields
is an independent author, speaker, and IT consultant. He's a vExpert, a Microsoft MVP, and a partner and principal technologist with Concentrated Technology. He serves as a technical guide for virtualization for *Windows IT Pro*.

Managing ABE from the Command Line

This PowerShell solution lets you control ABE on multiple shares

by Bill Stewart

Access-based enumeration (ABE) is a feature that first appeared in Windows Server 2003 SP1. In short, it prevents users from seeing files and directories they can't access when they navigate to a share. When ABE is disabled, users can see the existence of all files and folders when they open a share. If they try to open an item they don't have permission to access, they get an "access denied" message. In contrast, if ABE is enabled, the user simply doesn't see the files and folders they don't have permission to access.

ABE is disabled by default in Windows Server 2003 but enabled by default in Windows Server 2008 and later. You can change that default using the Abecmd.exe command-line tool. However, this tool is missing an important functionality, so I created an alternative solution.

Abecmd.exe's Downfall

Abecmd.exe is part of the Windows Server 2003 Access-Based Enumeration package at www.microsoft.com/downloads/details.aspx?FamilyId=04A563D9-78D9-4342-A485-B030AC442084. This Windows installer package also includes a white paper on ABE and an ABE GUI for Windows 2003. (Windows 2008 and later includes a built-in GUI that's part of the Share and Storage Management console.)

With Abecmd.exe, you can easily enable or disable ABE on a computer's shares on a case-by-case or global basis. However, you can't view the ABE status for multiple shares on a computer at the same time. I'm perplexed as to why this functionality wasn't included in the command-line tool. With this functionality, you could easily manage ABE on multiple shares because you could use a script to check each share's status and change it if needed.

So, I created a solution that lets you easily manage ABE on multiple shares and even multiple computers. This solution consists of three components:

- ShareABE.exe. This command-line program can detect whether ABE is enabled, enable ABE, and disable ABE on a single share. When detecting ABE, it returns an exit code of 1 if ABE is enabled or 0 if it's disabled. When you use it to enable or disable ABE, it returns an exit code of 0 to indicate success. Any other number indicates a failure. If a failure occurs, ShareABE.exe returns a real error code. For example, it returns an error code of 53 when the network path wasn't found, 2310 when the share didn't exist, and 5 when access was denied. You can use the Net Helpmsg command to check the error code.
- Get-ABE.ps1. This PowerShell script uses ShareABE.exe to detect whether ABE is enabled for one or more shares on one or more computers.
- Set-ABE.ps1. This PowerShell script uses ShareABE.exe to enable or disable ABE for one or more shares on one or more computers.

You must be a member of the Administrators group to change the ABE state. If you're changing the ABE state for shares on the local computer, you must launch PowerShell or Cmd.exe under elevated

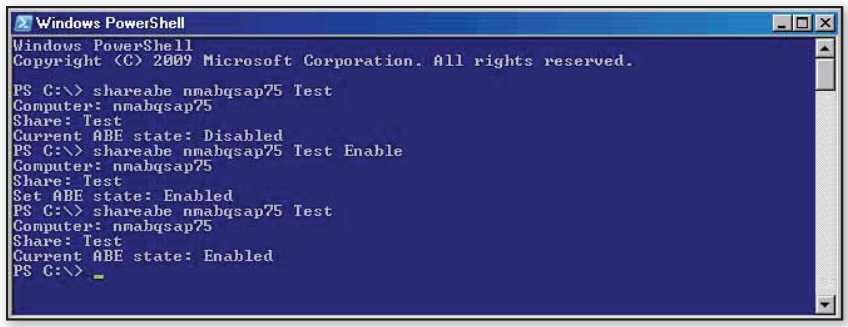

Figure 1: Sample output from ShareABE.exe when run as a standalone utility

permissions. To change ABE on a share for a remote computer, you must be a member of the Administrators group on the remote computer; elevation doesn't matter in this case. If you aren't a member of the Administrators group on the computer hosting the share, ShareABE.exe returns an error code of 5.

You can download this solution by going to www.windowsitpro.com, entering 129552 in the InstantDoc ID box, clicking Go, then clicking the *Download the Code Here* button. Besides including the two scripts and the command-line program, the 129552 .zip file includes the program's Free Pascal source code in case you're curious.

The solution works on Server 2003 and later, and unlike Abecmd.exe, it even works on Windows 7. I have used the solution with NTFS shares. If you want to use ABE on DFS shares, see "How to implement Windows Server 2003 Access-based Enumeration in a DFS environment" (support.microsoft.com/kb/907458) or "How to enable Access-based Enumeration for a Distributed File System (DFS) share in Windows Server 2008 (support.microsoft.com/kb/961658).

To use the solution, copy ShareABE.exe, Get-ABE.ps1, and Set-ABE.ps1 to a directory in your Path (e.g., SystemRoot\system32). The scripts require PowerShell 2.0. If necessary, adjust the PowerShell script policy to enable scripts to run. I recommend using the RemoteSigned policy. (If you're unfamiliar with the PowerShell policies, see "Running PowerShell Scripts Is as Easy as 1-2-3," March 2010, InstantDoc ID 103427.)

ShareABE.exe

ShareABE.exe is the tool that makes the PowerShell scripts work. It's also useful as a standalone utility. To run it as a standalone utility, go to a Cmd.exe or PowerShell prompt and run the command

```
ShareABE ComputerName ShareName
  [Action]
```

where *ComputerName* is the name of the computer that hosts the share, *ShareName* is the share's name, and *Action* is either the word *enable* (to enable ABE) or *disable* (to disable it). If you omit the *Action* parameter, ShareABE.exe outputs whether ABE is enabled or disabled on the share.

Figure 1 shows ShareABE.exe in action. The first command shows that ABE is disabled for the share, and the second command enables ABE. The third command, which is the same as the first, shows that ABE is indeed enabled for the share.

Get-ABE.ps1

Get-ABE.ps1 turns ShareABE.exe into a tool that can report the ABE status for multiple shares on multiple computers. Get-ABE.ps1's command-line syntax is

```
Get-ABE [-ComputerName String]
  [-ShareName String[]]
```

(Although this command wraps here, you'd enter it all on one line in the PowerShell console. The same holds true for the other commands that wrap.) Both of the script's parameters, -ComputerName and -ShareName, are optional. If you omit -ComputerName, Get-ABE.ps1 assumes you want to check the shares on the local computer only. If you omit -ShareName, Get-ABE.ps1 assumes you want to check all the shares. The share name can contain wildcards. Here are some sample commands:

```
Get-ABE app1
```

This command reports on the ABE status of all shares on the server named app1. As this command and the following one show, including the parameters' names (-ComputerName and -ShareName) is optional.

```
Get-ABE app1,app2,app3 Scan*
```

This command reports the ABE status for those shares whose names start with the word *Scan* on the app1, app2, and app3 servers.

Get-ABE.ps1 also accepts pipeline input for the -ComputerName parameter. For example, the command

```
Get-Content Servers.txt |
  Get-ABE -ShareName Shares,Users
```

reports the ABE status of the shares named Shares and Users for each of the servers listed in the file Servers.txt (assuming one computer name per line).

Figure 2: Sample output from Get-ABE.ps1

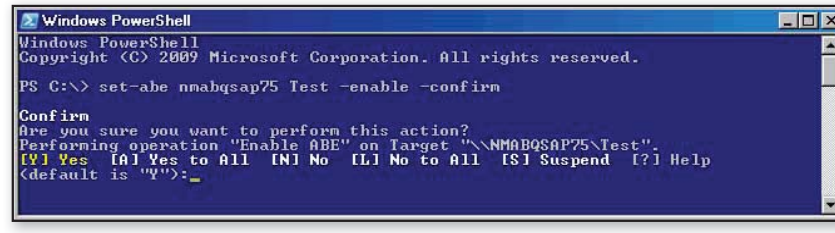

Figure 3: Sample output from Set-ABE.ps1

Get-ABE.ps1's output consists of custom objects containing the computer name, the share name, the share's path, and a Boolean value (True or False) indicating whether ABE is enabled for the share. Figure 2 shows a sample of Get-ABE.ps1's output.

Set-ABE.ps1

Set-ABE.ps1 is similar to Get-ABE.ps1, except that it sets the ABE state for shares instead. Its command-line syntax is

```
Set-ABE [-ComputerName String[]]
  -ShareName String[] [-Enable]
  [-Disable] [-WhatIf] [-Confirm]
```

The -ComputerName parameter is optional. If you omit it, Set-ABE.ps1 assumes the local computer. The -ShareName parameter is required and supports wildcards. So, for example, if you want to enable or disable ABE for all shares, you'd use an asterisk (*) to match all share names. You must include either the -Enable or -Disable parameter (but not both). For testing purposes, Set-ABE.ps1 also supports the -WhatIf and -Confirm parameters. Here are some sample commands:

```
Set-ABE apps1 * -Enable
```

This command enables ABE for all shares on apps1. This example omits the parameters' names (-ComputerName and -ShareName) because both parameters are present on the command line in the order specified by the syntax.

```
Set-ABE apps1,apps2,apps3 Users -Disable
```

This command disables ABE for the Users share on the apps1, apps2, and apps3 servers.

Like Get-ABE.ps1, Set-ABE.ps1 supports pipeline input for computer names. For example, the command

```
Get-Content Servers.txt |
  Set-ABE -ShareName Dep* -Enable
```

enables ABE for those shares starting with *Dep* on each server listed in the Servers.txt file (assuming one computer name per line). Figure 3 shows Set-ABE.ps1 in action, including the -Confirm parameter.

Take Control of ABE

ABE is a helpful tool in the administrator's toolkit, but it lacks a flexible, scriptable administration interface. ShareABE.exe, Get-ABE.ps1, and Set-ABE.ps1 rectify this shortcoming and make it possible to detect, enable, and disable ABE on multiple shares and on multiple computers.

InstantDoc ID 129552

Bill Stewart
(bstewart@iname.com) is a scripting guru who works in the IT infrastructure group at Emcore in Albuquerque, New Mexico. He has written numerous articles about Windows scripting and is a moderator for Microsoft's Scripting Guys forum

FEATURE

Exchange Server 2010 Namespace Planning

A key to setting up Exchange Server

by Siegfried Jagott and Joel Stidley

Editor's Note: *This article is excerpted from* Microsoft Exchange Server 2010 Best Practices *(Siegfried Jagott and Joel Stidley, Microsoft Press, 2010) and is reprinted with the publisher's permission.*

Before you set up your Microsoft Exchange Server organization, one of the most important areas that needs to be planned is your internal (organization-facing) and external (Internet-facing) namespace. A namespace is a logical structure commonly represented by one or more domain names in DNS.

Namespace planning is most important for the Client Access server role. However, many considerations are also needed for the Hub Transport and Edge Transport roles. This article provides the general basis for understanding the importance for namespace planning. The official Microsoft support statement for Exchange 2010 and Single Label Domain (SLD)/Disjoint/Non-contiguous Namespaces can be found at msexchangeteam.com/archive/2009/10/27/452969.aspx.

Namespace Scenarios

When you implement your Exchange Server 2010 organization, you need to decide how your internal and external namespace will be defined. This is important because it affects the following areas:

- DNS configuration of your Exchange servers
- How your certificates are created and what names they include
- Client Access (Outlook Anywhere—OA, Outlook Web App—OWA, POP3 and IMAP4, SMTP)

If you have multiple data centers available where your Exchange 2010 servers are located, consider the following general advice for namespace planning:

- Plan your namespaces such that both data centers can be active. This still allows for incremental deployment. You provide failover capabilities or can manually switch over a data center.
- Each data center needs the following namespaces, depending on your client connectivity capabilities: Outlook Web App/OA/Exchange Web Services (EWS)/Exchange ActiveSync (EAS) namespace, POP3/IMAP4 namespace, RPC Client Access namespace, and SMTP namespace
- Consider which data center will maintain the Autodiscover namespace.

To start planning your namespace, you need to consider the various locations of clients and servers and the physical connections they have to the Exchange servers. Typically the namespaces align with your DNS configuration.

You can choose from the following namespace-planning options:

- Consolidated data center
- Single namespace with proxy sites
- Single namespace with multiple sites
- Regional namespaces
- Multiple forests

Consolidated Data Center

This namespace scenario is the simplest one and includes a single namespace to access a single physical site where all the Exchange servers are hosted. This scenario has the following advantages:

- Only one or very few DNS records need to be managed.
- Only one or very few certificates are required for your Exchange organization.
- All users use the same URL to access the Exchange server.

This namespace scenario is configured by providing Internet access to the Client Access server by opening the relevant ports by a firewall or implementing an application layer firewall such as Microsoft Forefront Threat Management Gateway (TMG) in the perimeter network.

If you want to provide POP3/IMAP4, you also need to consider how the clients will send their messages using SMTP. To overcome this easily, you can configure the Hub Transport role on each Client Access server. Otherwise, you need to plan separately for message sending and message receiving namespaces.

Single Namespace with Proxy Sites

This model is based on the consolidated data center model but proxies the requests to the physical mailbox server located at another site. One of the sites has one or more Internet-facing Client Access servers that proxy the requests.

This scenario has the following advantages:

- Only one or very few DNS records need to be managed.
- Only one or very few certificates are required for your Exchange organization.
- All users use the same URL to access Exchange server.

The disadvantage of this model is that most users will access their mailboxes using proxying, thus accessing their data might be slower across latent WAN links.

To configure this namespace model, you need to configure the ExternalURL option of the Client Access server(s) at one site, and make sure that the ExternalURL settings on all the other sites are configured to $Null. This configuration ensures that the Client Access server does not redirect the connection to the target Client Access server, but instead proxies it. Redirect means that the Client Access server forwards the connection to the target Client Access server; proxy means that the Client Access server contacts the target Client Access server and retrieves the data for the connection.

Single Namespace with Multiple Sites

This model uses a single namespace for an organization that has multiple sites. A possible candidate for this approach would be a company that has multiple physical sites and wants to use a single namespace. The two possible approaches to implementing a single namespace with multiple sites are with a Client Access server proxy site or an intelligent firewall:

- The Client Access server proxy site approach includes Client Access servers based in a separate Active Directory (AD) site that is used to proxy the traffic to the site where the user's mailbox is located. To configure this namespace model, you need to configure the ExternalURL option to the single namespace of the Client Access servers at all sites.
- The intelligent firewall approach uses an application-layer firewall such as Forefront TMG and decides during client authentication that the traffic is forwarded to the correct target site based on configured rules. To configure this namespace model, you need to configure the ExternalURL option to the single namespace of the Client Access servers at all sites.

This scenario has the following advantages:

- Only one or very few DNS records need to be managed.
- Only one or very few certificates are required for your Exchange organization.
- All users use the same URL to access Exchange server.

The disadvantage of this model is that you must either have an application-layer firewall that is capable of forwarding the traffic to the correct physical sites, such as Forefront TMG or a Client Access server proxy site.

Regional Namespaces

This model uses one namespace for each region or site. The users will use their regional namespace to access their messages.

This scenario has the following advantages:

- The client traffic is automatically optimized based on the region or site level. (For example, if you implement a namespace based on a city, all users of that city will use the local access.)
- Performance and end-user experience are optimized.
- Failover is provided if the regional namespace is unavailable by using a different namespace (if the mailbox server of the site is still available).

The disadvantage of this model is that you need to manage multiple DNS records as well as multiple certificates. In addition, you have multiple Internet entry points that require a firewall.

Note: The regional namespaces model is recommended if your topology includes multiple sites that have their own Internet connectivity.

Multiple Forests

The multiple forest model uses one dedicated namespace for each forest. For example, if Contoso and Litware merged, Contoso users would need to access mail.contoso.com and Litware users would use mail.litware.com to access their mailboxes. Client Access server proxy redirection between the two forests does not work, so if one forest is not available, no users would be able to access their messages.

In this model, every namespace that is implemented needs its own Internet access point, DNS record, and a certificate. Within each forest, use a regional namespace model to improve customer experience.

Disjoint Namespace

The disjoint namespace model is a special scenario for planning the namespace. You face this scenario when your primary DNS suffix on domain controllers (DCs) or member servers in the domain is not the

same as the DNS domain name of your AD domain.

For example, you have a disjoint namespace when the Exchange server that is part of the Litware.com domain has a primary DNS suffix of Contoso.com. This computer (as the primary DNS suffix that does not match the DNS domain name) is said to be disjoint.

You might require these namespaces to be different for several reasons. For example, if DNS management in your company is split between administrators who manage AD and administrators who manage networks, you might need to have a topology with a disjoint namespace.

Exchange Server 2010 has three supported scenarios for deploying Exchange in a domain that has a disjoint namespace:

- Scenario 1: The primary DNS suffix of the DC is not the same as the DNS domain name. Computers that are members of the domain can be either disjoint or not disjoint.
- Scenario 2: The Exchange servers in an AD domain are disjoint, even though the DC is not disjoint.
- Scenario 3: The NetBIOS domain name of the DC is not the same as the subdomain of the DNS domain name of that DC.

In Exchange Server 2010 you might need to configure the DNS suffix search list to include multiple DNS suffixes if you have a disjoint namespace.

In a disjoint namespace environment, you must configure the following:

- All disjoint domains need to be added to the msds-allowedDNSSuffixes attribute of your root domain.
- The DNS suffix search list must include all DNS suffixes, including the disjoint DNS suffixes.

As mentioned, it is required to configure every disjoint domain in the msds-allowedDNSSuffixes attribute of your root domain. For example, if you have the disjoint namespace contoso.com that you need to add to the Litware.com forest, configure the settings on the domain level using the Windows Server 2008 tool ADSI Edit, as Figure 1 shows.

In addition, make sure that the DNS suffix search list contains all DNS namespaces that are deployed within your organization. To do this, you must configure the DNS search list for each computer in the domain that is disjoint. The list of namespaces should include not only the primary DNS suffix of the disjoint member computer and the DNS domain name, but also any additional namespaces for other servers the Exchange servers might interoperate with (such as the monitoring server). For more information on Exchange Server 2010 and disjoint namespaces, see "Understanding Disjoint Namespace Scenarios" at technet.microsoft.com/en-us/library/bb676377.aspx.

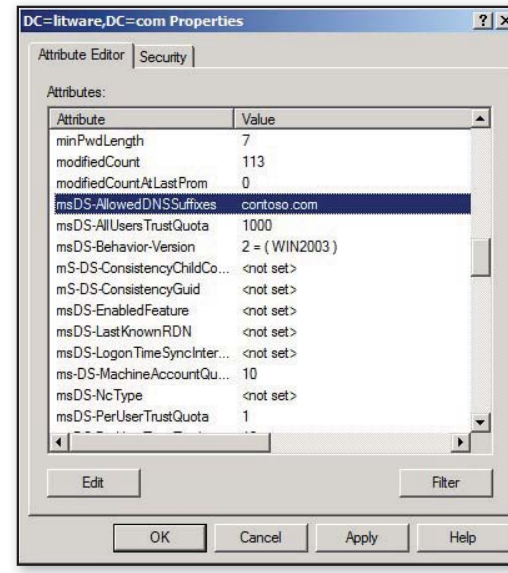

Figure 1: Configuring a disjoint domain

Single Label Domains

A single label domain (SLD) is basically a DNS domain name set equal to a NetBIOS domain name. It does not contain a suffix such as .com or .org and consists only of a single word, such as LITWARE or CONTOSO.

Before AD, in Windows NT an SLD was the standard, so some companies continued to use SLDs in AD. In Windows Server 2008 R2 you can no longer create SLDs—if you find an environment that still has SLDs, consider migrating to a normal namespace to prevent issues in the future.

Exchange Server 2010 supports SLDs; however, the Exchange product team does not recommend this configuration because future versions of Exchange or third-party applications might cause issues in this scenario. For that reason, you should move your organization to a normal namespace scenario.

Noncontiguous Namespaces

A noncontiguous namespace (sometimes referred as a discontiguous namespace) is a namespace in which an AD forest includes multiple domain trees of different names. Thus the forest is not defined hierarchically. A forest can have one or more domain trees, and these trees are defined by the DNS names. For example, contoso.com would be a domain tree in the Litware.com forest.

In Server 2008 R2, you can configure multiple domain trees by using Advanced Mode Installation in the Active Directory Domain Services Installation Wizard (dcpromo.exe). If you have similar tree names (such as litware.com and litware.de), be sure to choose different NetBIOS domain names for their respective domains. If you select the same NetBIOS names for both trees, the configuration is not supported. The general rule is that each domain must still register a unique legacy NetBIOS domain name.

If your organization has a noncontiguous namespace scenario, DNS must be configured so that every Exchange server is able to resolve all domain names in the environment. You are also required to configure msDS-AllowedDNSSuffixes within the AD environment for all namespaces used in the forest.

InstantDoc ID 129501

Siegfried Jagott
(siegfried.jagott@siemens.com) is a topic manager and senior architect at Siemens in Germany. He's an MCSE and presents at conferences on Windows, messaging, and collaboration topics. He's author of *Microsoft Exchange Server 2010 Best Practices* (Microsoft Press).

Joel Stidley
(joel@mailtask.com) is a hosting technology specialist at Microsoft and a former Microsoft Exchange MVP. He is the coauthor of *Microsoft Exchange Server 2010 Best Practices* (Microsoft Press) and maintains an Exchange community at exchangeexchange.com.

“THE CONVERSATION BEGINS HERE”

UNIFIED COMMUNICATIONS CONNECTIONS

Microsoft® Exchange CONNECTIONS

WINDOWS CONNECTIONS

SharePoint CONNECTIONS

Microsoft® Visual Studio® CONNECTIONS

Microsoft® ASP.net CONNECTIONS

Microsoft® Silverlight® CONNECTIONS

SQL Server CONNECTIONS

QUESTIONS ANSWERED • STRATEGY DEFINED • RELATIONSHIPS BUILT

NOVEMBER 1-4, 2011
LAS VEGAS, NV
MANDALAY BAY RESORT & CASINO

Make **CONNECTIONS** the **CONFERENCE** you bring your whole team to this year!

Only Microsoft and industry experts speak at WinConnections!

WinConnections ...
Providing the **vision**
+ intelligence
to keep you
and your company
competitive
in today's market!

QUENTIN CLARK
MICROSOFT

STEVE FOX
MICROSOFT

SCOTT GUTHRIE
MICROSOFT

MARK MINASI
MINASI RESEARCH AND DEVELOPMENT

JIM MCBEE
ITHICOS SOLUTIONS

KEVIN LAAHS
HP

KIERAN MCCORRY
HP

MIKE DANSEGLIO
CONCENTRATED TECHNOLOGY

DON JONES
CONCENTRATED TECHNOLOGY

CHECK WEB SITE FOR DESCRIPTIONS OF SESSIONS AND WORKSHOPS
www.WinConnections.com • 800.505.1201 • 203.400.6121 • Register Today!

Microsoft® SharePointPro CONNECTIONS SQL Server Magazine Windows IT Pro TECH Conferences Inc. PENTON MEDIA

SharePoint 2010 EASY as 1 2 3 !

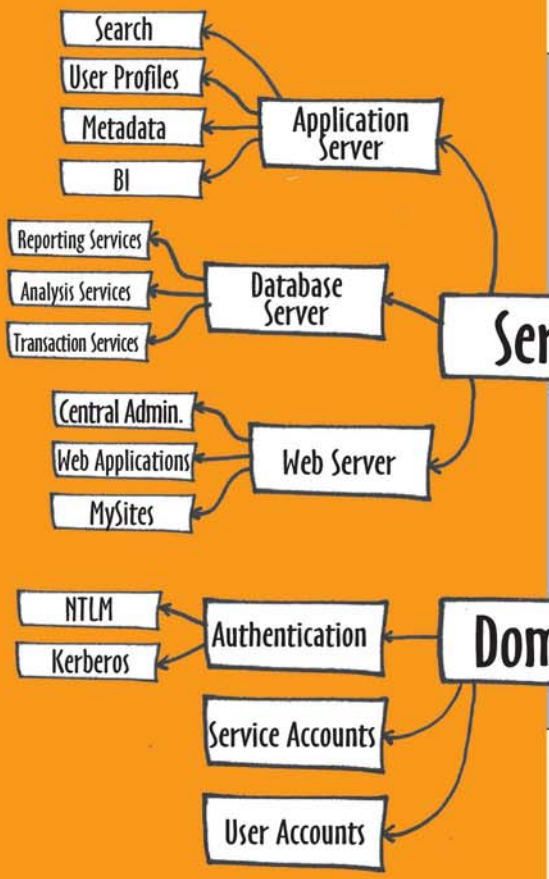

with SharePoint Composer® & Maestro®

- **Leverage Best Practices** from SharePoint MVP's & Certified Masters
- **Audit SharePoint** for Compliance with Governance Policies
- **Document & Save** existing configurations in Microsoft Word®
- **Detect invalid configurations**
- **Build SharePoint 2007 & 2010 Farms** for various SharePoint editions
- **Export configurations** to XML
- **Configure** SharePoint servers automatically from an XML configuration file
- **Import configurations** from existing 2007 farms and redesign for SharePoint 2010

1 Document

Crawl and import existing SharePoint farms to produce a rich visual depiction of your farm topology. Instantly document the configuration and detailed settings in Microsoft Word® to monitor ongoing usage and changes.

2 Design

Turn a deeply technical exercise into a visual, intuitive experience that significantly eases the design and configuration of SharePoint, from the farm level down to the site level.

3 Deploy

Once you've fine-tuned the design of your SharePoint farm, deploy it across multiple environments without writing complex scripts or performing tedious, error-prone manual configuration.

Download a FREE TRIAL at SharePointComposer.com

Share squared

Contact Us: 800.445.1279
Info@ShareSquared.com

SharePoint 2010 Disaster Recovery

Recover data lost to accidental content deletion

by Todd O. Klindt

I don't know about you, but I think of my SharePoint servers almost like children. I've been known to brag them up before. I love hearing about when they do good things, and when someone says something bad about them I always leap to their defense, whether they deserve it or not. And as with my children, I don't want bad things to happen to them. Fortunately it's much easier to make SharePoint servers behave than it is children, and SharePoint servers seem to bounce back more easily too. But no matter how hard we try to protect our SharePoint servers (and our kids), disasters do happen.

This is the first of two articles in which I cover the types of disasters that might befall your SharePoint servers and discuss ways to protect them from the cruel world around them. I also discuss ways to recover from disasters, whether you've prepared for them or not. Although a disaster might strike your SharePoint farm at some point, with the information in these articles you'll be able to confidently recover your data and restore your farm to its former glory in no time. In this article I cover disasters related to content deletion. I show different ways to prepare for this disaster, including ways to respond if it happens. In the next article, I'll cover disasters that revolve around hardware failure, including SharePoint servers, SQL Server machines, and your actual facility.

Before we get too far into our talk about disaster recovery, we need to agree on what a disaster is. Different types of disasters can affect your SharePoint servers. To be prepared for a disaster, you need to know what type of disaster you're preparing for. Of course in a perfect world you'd protect your SharePoint servers against all disasters. But unless you have one of those geese that lays golden eggs, you probably don't have the resources to take all the steps I'll cover. You'll most likely need to look at your budget and see which disasters you'll be able to protect your servers against.

Determining Which Disasters to Plan For

Before we cover how to deal with the disasters that might be headed your way, we need to discuss how to determine which of them you want to protect against. Of course, we all want to protect our SharePoint farms from all attacks, but because of budgetary and time constraints that's not always possible. If it's not possible in your environment, you need to give some thought to which disasters you want to protect against, so that you know what steps to take to protect yourself against them. The best way to do this is to discuss it with your IT management. Lay out the types of disasters, the processes for protecting against them, and the rough cost of each option. This will give you an idea where each option sits within your budget. Maybe you have more or less money than you thought you did for this project. Knowing the numbers helps.

Next, you should talk to your customers, the business units. Outline the different options with them and explain the costs. Maybe they're okay with not being able to do point-in-time recoveries, but maybe every minute of downtime costs them sales, which costs them money. Having the discussion

with them ensures that you understand how they use SharePoint and what's important to them. It also helps them understand the costs and time associated with each disaster. If they must have five 9s of uptime, or point-in-time recoveries, then they'll know the associated costs and you'll be able to work together to make sure the money is found to achieve their goal.

After you determine which disasters to protect against, the terms need to be spelled out in a service level agreement. The SLA outlines many operational expectations, and disaster recovery is one of them. This agreement outlines the terms of service. For example, suppose someone demands that a file deleted at 3:30 A.M. must be recovered with all its contents. If the SLA says that the only files that can be recovered are ones that existed at midnight, then you have something to support you when you tell the customer you can't recover the file. The SLA sets everyone's expectations the same and can prevent hard feelings if IT can't do something a customer wants done.

Recovering Lost Content with the Recycle Bin

Probably the most frequent disaster in the SharePoint world is prematurely deleted content. This might not seem like a disaster-level problem, but trust me, it is—especially to the person who erroneously deleted the content. And depending on this person's location on the company's organizational chart, a disaster for him or her might also mean a disaster for you.

The good news is that although this is the most common disaster you'll encounter, it's also the easiest to deal with. Like SharePoint 2007 before it, SharePoint 2010 has a built-in Recycle Bin. This Recycle Bin captures deleted list items, documents, folders, lists, and document libraries. Although this seems pretty comprehensive, keep in mind that the Recycle Bin doesn't catch web or site collection deletions (which I cover later). The Recycle Bin has two parts: the end user Recycle Bin and the administrator Recycle Bin.

The first stage of the Recycle Bin shows a user all the documents he or she has deleted in the current web and gives the user the option of restoring those documents. Figure 1 shows how a site member sees the Recycle Bin. The Recycle Bin lists the items the user has deleted in a site, including the items' original locations and when they were deleted.

Probably the most frequent disaster in the SharePoint world is prematurely deleted content.

From here, a user can restore a document to its original location and get back to work on that important document. Although the documents are in the first-stage Recycle Bin, they count against the site collection's quota. Therefore, end users can also delete items from the first-stage Recycle Bin to free up space. At that point, users can no longer restore items and must call the Help desk—which is where the second stage of the Recycle Bin comes in.

Unbeknownst to the end users, there's a second, secret Recycle Bin that's exposed only to site collection administrators. This Recycle Bin contains all the documents that every user has deleted in the site collection—and more important, it also contains all the documents that have been deleted from the first-stage Recycle Bin. This is where the SharePoint administrator can save users from themselves. If a user deletes a document and also deletes it from the first-stage Recycle Bin, it can still be restored from the second-stage or administrator Recycle Bin. Items in the administrator Recycle Bin also don't count against the site collection's quota. Those Microsoft folks thought of everything.

You need to be a site collection administrator to access the administrator Recycle Bin. To find it, open the first-stage Recycle Bin. Depending on which site template you're using, there's typically a link on the left navigation bar. If you're a site administrator, you'll see additional text that mentions the Site Collection Recycle Bin, with a link to it. When you click that link you'll be taken to the /_layouts/AdminRecycleBin.aspx page. This page has two links. One shows the contents of the first-stage, or end user, Recycle Bins. The second link, called *Deleted from end user Recycle Bin,* shows the second stage. Here you can see all the documents end users have deleted from the Recycle Bin. Restoring a document from here will restore it to its original location.

It's important to recognize that deleted items don't age out of the first stage to the second stage. They only show up in the second stage if they were deleted out of the first. If the web application settings are

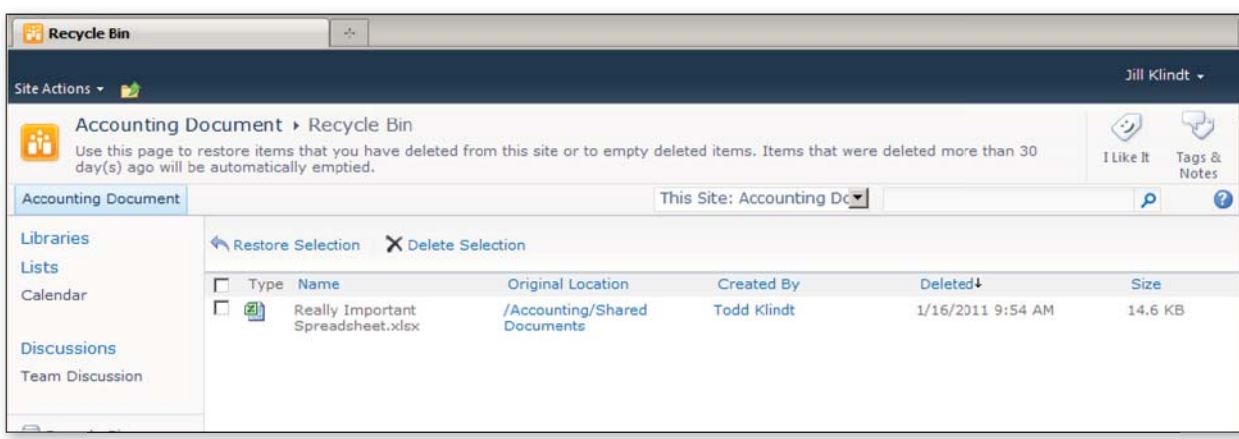

Figure 1: SharePoint 2010's Recycle Bin

set to delete items from the Recycle Bin after 30 days, then they are gone for good out of both stages of the Recycle Bin. There are no tricks to get them back.

This setting and the rest of the Recycle Bin settings can be accessed by selecting Central Admin, Application Management, Managed Web Applications. Choose a web application and click General Settings in the ribbon. If you want to flex your PowerShell muscle, you can also modify these settings with the RecycleBinCleanupEnabled, RecycleBinEnabled, RecycleBinRetentionPeriod, and SecondStageRecycleBinQuota properties of your web application.

Other Tools for Recovering Content

How do you recover content after it leaves your Recycle Bin? That's a little trickier. To accomplish this task, you need some sort of backups. If you're not doing any type of backups right now, first, shame on you. Second, you're lucky you're reading this article. The very least you can do is back up all your SharePoint databases in SQL Server. Armed with those backups, you have several options for restoring content. If you're not familiar with how to do database backups in SQL Server, you can get a quick primer from my blog article "Scheduling SQL backups for SharePoint" at www.toddklindt.com/blog/Lists/Posts/Post.aspx?ID=248.

You have a few options if you're recovering content from databases. Which option you choose depends on your environment and what you're trying to recover. If you need to recover a deleted site collection or Web, it doesn't get any easier than the Central Administration unattached content database recovery option. You can see in Figure 2 that Central Administration has an entire tab dedicated to backup and restore. One of the options lets you access data in a SharePoint database that isn't mounted in SharePoint. You can use this option to back up data from a recovered database and restore it to your production environment. To do this, restore an old backup of a content database into SQL Server. Use a different name to make sure you don't overwrite the production version of the database when you're restoring your backup database. Then click *Recover data from an unattached content database* in Central Administration. Enter the name of the recovered database, select Browse Content, and click Next.

> Thanks to the awesomeness that is PowerShell, you can easily back up all the site collections in your farm with a single line.

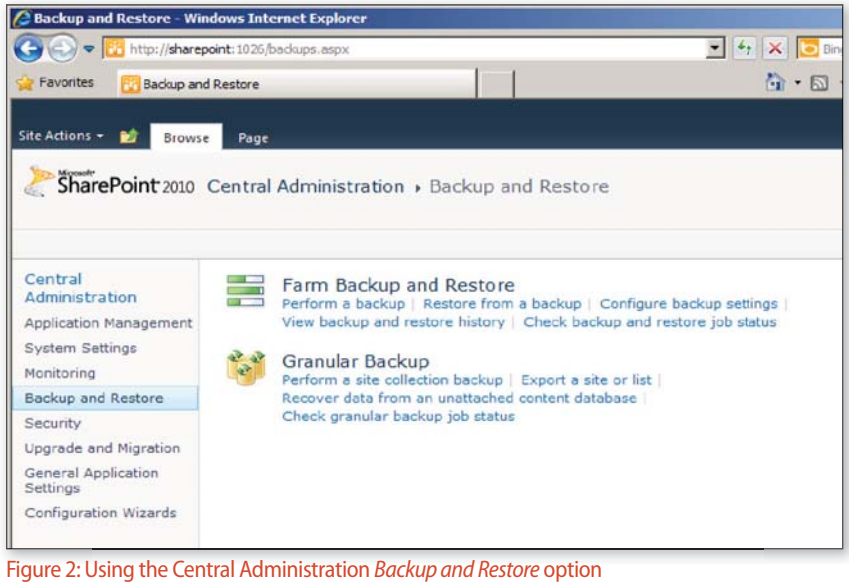

Figure 2: Using the Central Administration *Backup and Restore* option

From here, you can browse to a site collection that exists in the restored content database. You can also export just a web or list if you prefer. If you choose a site collection, the site collection is backed up as with Backup-SPSite or stsadm -o backup. If you choose a Web or list, the content is exported as with Export-SPWeb or stsadm -o export. When you click Next, you'll see a screen asking for a filename to save the backup or export file. It's important that you use a Universal Naming Convention (UNC), because the backup or export is done as a timer job and can run on any machine in your farm.

When you click Start Backup, SharePoint will begin backing up or exporting the content you selected to the appropriate file. A backup job status page then opens to let you know how your backup is going. This page also has a Recovery Step column for recovering content. For site collections, use Restore-SPSite. For webs and lists, use Import-SPWeb. When your backup finishes, use the correct command to restore your recovered content to either its original location or another location if you want to pull out individual items. The other location can be on the same farm or a separate farm. The only requirement is that the farm you recover to must be the same build number or later than the farm the backup is from.

If you have a second SharePoint 2010 farm, you can skip the unattached database steps and attach your restored content database to your second farm. The recovery farm must be at the same SharePoint build or later than the farm that created the database. If any solutions are needed to render the content, you'll likely need that solution on the recovery farm as well.

Another often overlooked setting is managed paths. The web application you restore the database to must have any managed paths needed to render out the content you're trying to recover. After the database is attached to the recovery farm, you can use several techniques to obtain the desired content: You can perform backups or exports, you can download individual documents, and you can use Explorer View to obtain several documents at once.

I already touched on this topic briefly, but you can also use PowerShell to perform backups of individual site collections.

The cmdlet to back up site collections is Backup-SPSite. Go to the SharePoint Management Shell and enter

```
Get-Help Backup-SPSite
```

to obtain usage information. For basic backups, you need to provide only two things: the URL of the site collection being backed up and the name of the file to back it up to. The file that the Backup-SPSite cmdlets writes will be a full fidelity backup of that site collection. This means the backup will contain the content plus its metadata (e.g., alerts, security, workflows). Like the content databases, this backup is completely portable. You can restore it to a different location in your farm or to a different farm altogether.

Although you probably wouldn't use this method as part of a disaster recovery strategy, it's a good alternative for performing one-off backups. Thanks to the awesomeness that is PowerShell, you can easily back up all the site collections in your farm with a single line. Again, this isn't a good enterprise disaster recovery technique, but it might come in handy once in a while.

```
Get-SPSite | ForEach-Object{$FilePath
  = "d:\backups\" + $_.Url
  .Replace("http://","")
  .Replace("https://","")
  .Replace("/","-").Replace(":","-")
  + ".bak" ; Backup-SPSite -Identity
  $_.Url -Path $FilePath}
```

This command will back up all the site collections in your farm to the D:\Backups folder on the server you run it on. We use the Replace() method to strip out parts of the URL that don't work in filenames (e.g., colons, slashes). Make sure your drive has enough space for the files before you run this command. You'd use Restore-SPSite to turn those files back into site collections.

Third-Party Options

As we've seen, SharePoint has a pretty good set of tools out of the box to help you prepare for content-loss disasters and recover from them in the unlikely event that they do occur. However, there are also some good third-party solutions you should take a look at. These aren't full disaster recovery suites; I'll cover those in the next article. These are just some tools to augment the out-of-the-box tools I already discussed.

SharePoint Site Recycle Bin. When I covered the SharePoint Recycle Bin earlier, I mentioned that it doesn't protect against accidental web or site collection deletion. I then explained how to recover webs and sites from database backups, but this process can be a lot of work. Also, the restored web or site collection will only be as current as the last backup. Fortunately there's a free tool to fill that gap. You can download the SharePoint Site Recycle Bin from the SharePoint Governance and Manageability website at governance.codeplex.com. This is a free download that adds Recycle Bin–type functionality to SharePoint.

These techniques will help you restore your users' content, no matter how hard they work to delete it.

Before a web or site collection can be deleted, SharePoint backs it up to a location in the file system. Then when the Help desk calls start rolling in, you can simply restore the web or site collection from the backup. The only way this process could be easier is if SharePoint answered the phone, chatted up the customer, and restored the customer's content for you. (I hear that's coming in the next version.)

The setup for the SharePoint Site Recycle Bin is a little tricky, so make sure you read through the instructions before you start the installation. I also recommend trying it out in your test environment first to get a feel for the setup. Finally, make sure the location you save the backups to has enough space and is getting backed up itself (or not getting backed up—whichever fits your retention policies).

Idera's SQL virtual database. I covered several ways to use SQL Server database backups to recover deleted content. If you want to take your database recovery game up another notch, you might consider buying Idera's SQL virtual database. This software lets you mount a database backup in SQL Server without actually restoring the database, which saves time and drive space. Instead of restoring the database back to SQL Server, SQL virtual database attaches to the backup file and exposes the virtual database to SQL Server, giving you immediate access and not requiring the drive space for the restored MDF and LDF files.

After the virtual database is mounted in SQL Server, you can use any of the methods I discussed to recover your content. This approach lets you keep more backups around on disk and lets you get your users' content back to them more quickly. For more information about SQL virtual database, go to Idera's SQL toolbox website at www.idera.com/Products/SQL-toolbox/SQL-virtual-database.

AvePoint's DocAve Recovery Manager for SharePoint. We've discussed a couple of ways to recover documents from database backups. Although the process isn't bad, it's awfully manual. Wouldn't it be nice to have a tool that you could just point at a recovered database and select the document to recover? AvePoint has created just such a utility—and it's free!

Go to AvePoint's website at www.avepoint.com/docave-recovery-manager-for-sharepoint to download the DocAve Recovery Manager for SharePoint. As with any software, you should implement and try it out in a test environment first to work out the kinks before you unleash it on your production servers.

Wrapping It Up

A multitude of disasters could befall your SharePoint farm. To be a top-notch SharePoint administrator, you need to stay a step ahead of these disasters. In this article I discussed several ways to recover deleted content and get it back into your farm. These techniques will help you restore your users' content, no matter how hard they work to delete it. I'll take a look at some other types of disasters, such as machine and facility failure, in a follow-up article and discuss how to plan for them.

InstantDoc ID 129713

Todd O. Klindt
(todd@sharepoint911.com) is a consultant for SharePoint911 and a SharePoint MVP.

- Cloud Identity
- Storage
- Network Monitoring
- Virtualization

Norton Introduces NortonLive Ultimate Help Desk

Sick of being everyone's personal IT pro? Norton by Symantec has announced the **NortonLive Ultimate Help Desk** service, which provides support for consumer devices, including PCs, printers, mobile devices, and more. The service is available 24x7. It costs $19.99/month for an individual or $29.99/month for a family. Is it a waste of money, or a small price to pay to get your weekends back? To learn more, visit us.norton.com.

Sans Digital Enhances Network Attached Storage Devices

Sans Digital has released a new series of 64-bit NAS products. Included are two tower-based, 4-bay models and eight

rackmounted models, ranging from 4 bays to 50 bays. In addition to the switch to 64 bit, enhancements include a new GUI with a pull-down menu to simplify configuration. The new NAS/iSCSI systems support ESX Server, Hyper-V, and XenServer virtual machines servers; Windows Server 2008 clustering; and SPC-3 persistent targeting. Optional features include 10GBit networking support and server failover/mirroring. Note that all devices run on Linux. To learn more, visit www.sansdigital.com.

PRODUCT SPOTLIGHT

Symplified Introduces Cloud Identity and Access Management Platform

Symplified has announced **Symplified Suite,** a cloud identity and access management platform. Capabilities include cloud web access management, federated single sign-on, and user management/provisioning with identity management. (The product has three components that correspond with these capabilities: Symplified Access Manager, Symplified Identity Manager, and Symplified Sign-On.)

"The cloud, in public, private, or hybrid incarnations, is reshaping the way organizations use information technology and manage access to systems, applications, and data," said Eric Olden, founder and CEO of Symplified. "Symplified, unlike legacy identity management solutions, was designed from day one for the cloud and to operate in the cloud. This architecture has allowed us to build a feature complete suite that seamlessly supports hundreds of SaaS applications, Amazon EC2 and now private clouds—years ahead of anyone else. With our latest release, Symplified offers the only unified access management platform for the cloud. Now enterprises don't have to compromise capabilities when moving to the cloud."

Identity management was ranked as a top priority in a survey of CIOs. Symplified Suite offers a unique, cloud-based answer to identity management. Other features of the product include: accepts identities based on the Open ID standard; an app store of pre-integrated cloud applications; and the ability to add single sign-on to an application in minutes.

To learn more, visit www.symplified.com.

Clone Your PC to a New Drive

NTI Corporation has announced **Echo,** a program that lets you clone a PC's old drive to a new drive. NTI has created Echo for bundling with solid-state drives (SSD), hard disk drives (HDD), and hybrid drives. Echo will clone an entire drive, including all of its partitions, with all of the user's data, applications, and the OS to another drive. The new drive can be of different types and sizes as long as the data can fit within its capacity. Echo will shrink and grow partitions as needed to optimize the use of available space. To learn more, visit www.nticorp.com.

ProQueSys Releases Network and Security Monitoring Software

ProQueSys has released **ProQueSys FlowTraq**, a network monitoring, security, and forensics flow analyzer. FlowTraq is designed to complement and improve existing network operations. It helps IT administrators find data leaks in the network, investigate compromises, and monitor network usage such as bandwidth consumption, applications in use, and changes in behavior or network activity that may indicate a problem. The product complements full-packet capture solutions. The product also offers blacklist address alerting and policy development. To learn

NEW & IMPROVED

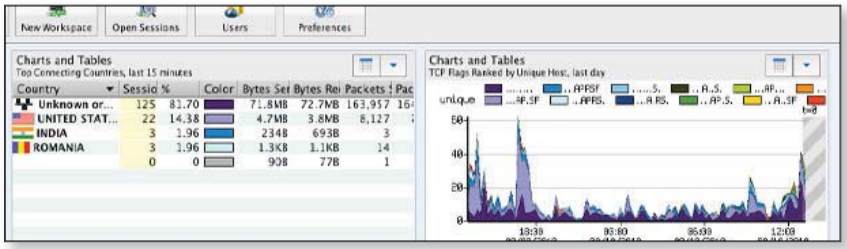

more, visit www.proquesys.com.

SPAMfighter Introduces FULL-DISKfighter

SPAMfighter today announced **FULL-DISKfighter**, a file and disk cleaning solution. The new software aims to help PC users quickly and easily delete unwanted junk files weighing down computer performance, improving performance and stability. The product has the following features: find unnecessary files; clean Windows Update history, temporary Internet files, service packs, log files, error reports, etc.; reorganize and defrag hard disks more efficiently; and locate large and duplicate files. To learn more, visit www.spamfighter.com.

F5 Accelerates VMware View Deployments

F5 Networks is making the **F5 BIG-IP Access Policy Manager (APM)** module licensable on the BIG-IP Local Traffic Manager Virtual Edition (LTM VE) product. The resulting product, shortened to **APM for LTM VE**, offers organizations an authentication, authorization, and accounting module for virtual application delivery controllers. Features include the ability to deploy VMware View quickly, control infrastructure costs, scale VMware View deployments, and simplify management through one central console. To learn more, visit www.f5.com.

Paul's Picks

www.winsupersite.com

SUMMARIES of in-depth product reviews on Paul Thurrott's SuperSite for Windows

Internet Explorer 9

PROS: Hardware acceleration, standards-based rendering, Windows 7 integration

CONS: Some features require Windows 7, lingering site compatibility issues

RATING: ◆◆◆◆◇

RECOMMENDATION: Microsoft's latest browser will face slow adoption in business, where website compatibility issues will bedevil admins and IT pros. But it's a no-brainer for individuals, with its superior UI, excellent standards-based web rendering, full hardware acceleration, and, on Windows 7, deep integration with the underlying OS. The only serious issues with IE 9 are that it won't run on Windows XP—still a huge percentage of the installed base—and that some of its best features—including website pinning—require Windows 7. But it's a viable alternative to Google Chrome and clearly preferable to Mozilla Firefox and Apple Safari.

CONTACT: Microsoft • www.microsoft.com

DISCUSSION: www.winsupersite.com/article/internet-explorer2/Internet-Explorer-9-Release-Candidate.aspx

Windows Home Server 2011

PROS: Solid Server 2008 R2 base, centralized PC and server backup, remote access, new add-in extensibility model

CONS: No more automatic data duplication and single pool of storage

RATING: ◆◆◆◆◇

RECOMMENDATION: When Microsoft removed Drive Extender from Windows Home Server 2011, enthusiasts protested. But WHS 2011 is still solid. And much of Drive Extender's functionality can be replaced by automated daily server backups and the server's underlying Previous Versions functionality. Combine this with WHS 2011's core features—centralized backup for connected PCs as well as the server, centralized media sharing, network health monitoring, remote access—and you get a product whose value is greater than the sum of its parts.

CONTACT: Microsoft • www.microsoft.com

DISCUSSION: www.winsupersite.com/article/windows-server/Windows-Home-Server-2011-Release-Candidate.aspx

InstantDoc ID 129762

Iomega StorCenter ix4-200d

NAS devices can be substantially cheaper than their data center–class counterparts, in part by sacrificing the use of high-end SCSI or Serial Attached SCSI (SAS) drives for Serial ATA (SATA) drives, and by using a proprietary version of Linux instead of Windows Storage Server. I considered these devices as I looked for a replacement for my aging and power-hungry DAS array, which is comprised of twelve 15,000rpm 72GB SCSI drives. My list of required features included support for common NAS file-serving protocols such as Server Message Block (SMB), Apple File Protocol (AFP), and NFS. I wanted a system that could integrate with my Active Directory (AD)-based network, as well as one that offered iSCSI support, so that I could rebuild my wholly virtualized test environment. I also wanted fault tolerance so that I wouldn't lose any data in the event of disk failure.

I finally settled on the Iomega StorCenter ix4-200d with 8TB of disk space. In addition to meeting my requirements, the system also comes with backup software for desktops and servers, provides remote access to files over the Internet, offers Blue tooth support, and has print server and USB storage expansion capabilities. I was drawn to Iomega's products because they're wholly owned by EMC, a heavyweight in SAN devices, and the ix4-200d is based on the same technology.

The ix4-200d is small, compact, and extremely well built. Two thumbscrews on the back let you access the four 2TB user-replaceable hard drives. On the front of the device is a display with two buttons that control what's shown. By default, the device displays its status and configuration information in a loop. There are USB ports on the front and back of the device, as well as two 1Gb Ethernet ports on the back. The Ethernet ports can be bonded for improved throughput or used on separate networks. Hooking up the ix4-200d to my network was as simple as plugging in an Ethernet cable and power.

The device comes with a CD-ROM that contains management and backup software, which was simple to install. When I launched the management software, there were a few glitches as it tried to find the ix4-200d; this process could be improved. After setup, you can dispense with the software and use a web browser to manage the device. I joined the device to my AD forest but then immediately put it back into workgroup mode because when it's joined to a domain, the support for AFP is disabled. Iomega should fix this problem, because many enterprise environments integrate Macs with AD. Joining and leaving a domain is a breeze.

Next, I reconfigured the device to use RAID 10 instead of RAID 5, which is the default. My primary use of the device is to serve Virtual Hard Disks (VHDs) to Windows Server 2008 running Hyper-V using iSCSI. RAID 10 is preferable to RAID 5 in this scenario because of the improved read/write performance. It took almost 24 hours to rebuild the volume, which seemed long.

Creating iSCSI volumes to serve was easy (as is creating ordinary file share access using SMB, AFS, or NFS). I created a user account and set a password to protect the iSCSI volumes. I also configured the device so that the second Ethernet port was on a private LAN used by my Hyper-V server. Despite the documentation clearly stating that you need to install the Iomega management software on Windows Server systems to access the iSCSI targets, I found that this wasn't true. The built-in iSCSI initiator worked just fine and let me specify a username and password to connect to the iSCSI targets. Although the process wasn't as simple as using Apple's Time Machine products, I was able to configure the device as a Time Machine backup server and both my Mac Mini and MacBook Air were able to connect to the ix4-200d.

Performance was stunning. I'm running 12 virtual machines (VMs), a mix of UNIX and Windows, whose hard drives are on the ix4-200d and are accessed using iSCSI. Although I wouldn't attempt this in a production environment, the device is perfect for my test lab—I'll probably end up running 20 or more VMs from the device. Performing installations of OSs does slow

the device down, but that might simply be because of extremely high network utilization rather than a bottleneck in the device itself. After installation, everything appeared to be just as fast as when using my DAS.

I bought the device for use as an iSCSI target server for my virtualized test environment, and the device is excellent for this purpose. However, I quickly pressed it into service as a Time Machine backup server for my Macs, as well as for my Windows desktops and servers.

The Iomega ix4-200d offers an impressive array of features. The device is useful for small businesses or for larger organizations with departments that are looking for flexible storage options, as well as for test labs and even some production environments hosting VMs.

InstantDoc ID 129566

Iomega StorCenter ix4-200d

PROS: Supports numerous enterprise file-sharing protocols; full iSCSI target support; stunning performance; impressive all-up feature list

CONS: Active Directory integration disables support for Apple protocols; clumsy management software; inaccurate documentation at times

RATING: ◆◆◆◆◇

PRICE: $1,300 for 8TB model

RECOMMENDATION: This extremely flexible and high-performing device will meet most administrators' backup, storage, and performance needs.

CONTACT: Iomega • 858-314-7000 • www.iomega.com

John Howie | john@thehowies.com

REVIEW

HP Business Decision Appliance

Most organizations understand the importance of implementing business intelligence (BI) solutions. BI and data analysis are vital tools for transforming the raw data that's generated by an organization's line-of-business (LOB) applications into intelligent information that can be used to drive the business decision-making process.

However, there are several hurdles to implementing BI solutions. BI technologies are different from the relational technologies that are used to support most business applications. Implementing performant and scalable BI solutions requires a different hardware platform and skill set from the relational databases that drive LOB applications. The HP Business Decision Appliance is designed to address these technological hurdles by providing a ready-to-run BI appliance that lets businesses quickly deploy BI solutions throughout the organization.

I reviewed the new HP Business Decision Appliance (BDA) at the Microsoft Enterprise Engineering Center (EEC) on the Microsoft campus in Redmond, Washington. The BDA is designed to be a BI platform for small and midsized businesses (SMBs), as well as branch offices or departments in an enterprise. The BDA provides more than adequate scalability for most SMBs; in an enterprise scenario, it can act in concert with a larger enterprise data warehouse, such as the Microsoft SQL Server 2008 R2 Parallel Data Warehouse, in a hub-and-spoke arrangement whereby the BDA functions as a data mart for different subsets of the organization.

Under the Hood

The HP BDA is based on the HP ProLiant DL360 G7. Unlike a standard server that you might buy from an OEM, the BDA comes preconfigured. There's nothing for the customer to worry about or change. The BDA is a 1U appliance that comes with dual Intel six-core X5650 2.67GHz Xeon processors. It's equipped with 96GB of RAM and eight 10,000rpm 300GB Serial Attached SCSI (SAS) drives. The first two drives are mirrored for redundancy and contain the system software. The remaining six drives are configured in a RAID 5 array that provides a total of 1.5TB of storage.

The RAID 5 array on the unit I tested was split into a 653GB D drive that contained the Microsoft SharePoint database and logs and a 713GB E drive that was used for system backup. Because the appliance is intended as a one-step install-and-run type of device, there are no PCI slots for additional add-on components.

The front panel of the appliance houses a VGA video port, one USB 2.0 port, and an HP Systems Insight Display (SID) that provides easy-to-access system diagnostic information. Although I wished there were more USB ports on the front, I did like the easy access to the system diagnostics.

The back of the unit has two USB 2.0 ports, four 1GB Ethernet ports, one VGA video port, one serial port, and one Integrated Lights-Out (iLO) port for remote management. Notably, like many of the newer server units, there are no PS/2-style mouse and keyboard ports. The unit also lacks a DVD or optical drive. The system has two 460W hot-swappable power supplies. The BDA is essentially designed with the goal of efficiently serving multiple PowerPivot workbooks using SharePoint.

Deploying the BDA

Installing the appliance was surprisingly easy, considering that the process essentially consisted of installing three different products: Windows Server 2008 R2 Enterprise Edition, SQL Server 2008 R2 Enterprise Edition, and SharePoint 2010 Enterprise Edition. To begin the installation, I first connected to the BDA using HP's iLO 3 management software. Like a standard HP server, the BDA supports full out-of-band management using iLO. iLO lets you power the appliance on and off, as well as perform other systems management tasks, such as running system diagnostics, checking system logs, and monitoring the status of the system's hardware components. Using iLO on the BDA was exactly like using it on a standard HP server. Figure 1 shows the iLO console connected to the BDA.

After the appliance initially powered, I was presented with a series of setup dialog boxes. The initial dialog box prompted me to accept the EULA information. Next, I was presented with a dialog box asking me to change the appliance's initial administrator password. After I logged on to the appliance as an administrator, the setup process displayed the HP Business Decision Appliance Quick Deployment Tool dialog box that Figure 2 shows.

The HP Business Decision Appliance Quick Deployment Tool guided me through the process of installing the BDA. To complete the setup, I also referred to the HP Business Decision Appliance Installation Overview guide, which provided clear and easy-to-follow, step-by-step instructions for setting up the appliance. The initial dialog box states the installation requirements.

To complete the installation of the BDA, I needed a connection to an Active Directory (AD) domain, a connection to a DHCP server, physical access to the appliance or access to the iLO connection, a domain user account to add the appliance to the domain, and a domain service account to run the SQL Server and SharePoint services. A handy More Details button on the HP Business Decision Appliance Quick Deployment Tool screen provides additional explanation about the installation requirements.

When I clicked Next, I was presented with a *Machine Name and Domain Join*

Michael Otey | motey@windowsitpro.com

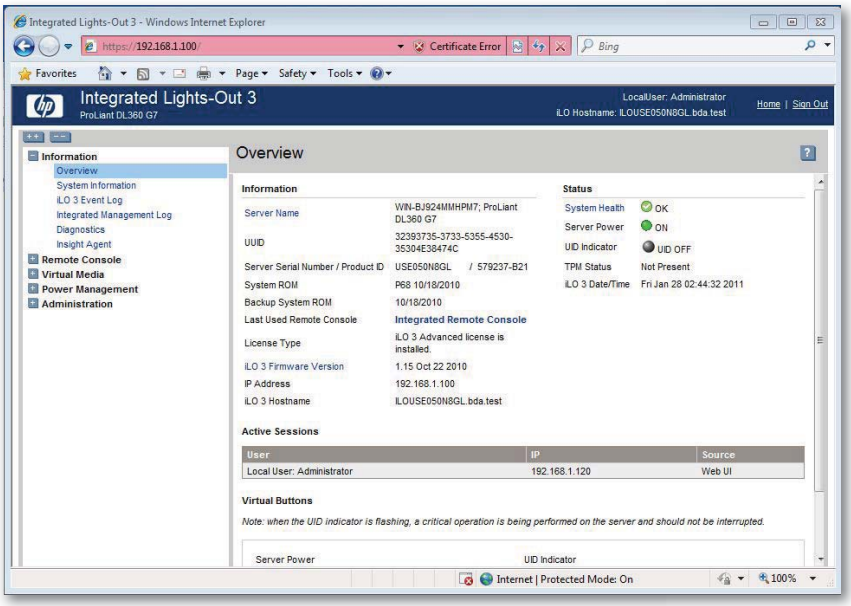

Figure 1: Integrated Lights-Out management

dialog box. This dialog box prompted me to input the appliance name, the domain name, a domain username that has rights to add a computer to the domain, and finally the password for the domain account.

It's important to realize that unlike a typical SQL Server or SharePoint installation, the installation of the BDA actually adds a new computer to the domain. Therefore you need domain rights to add the new computer.

I named the appliance BDA01, provided the authentication information for the domain in the Microsoft EEC, and clicked Next. After a brief pause, the machine rebooted and presented another EULA for both HP and Microsoft agreements. Again, I entered the domain and authentication information. The setup applied various roles to the appliance, including the IIS Web Server and Application Server roles, then rebooted.

The final stage of the installation requested that I supply the domain service account for the SQL Server and SharePoint services. It also asked for a SharePoint farm security passphrase and the domain user accounts that can act as database administrators. The domain administrator was added by default to the list of SQL Server administrators. Clicking Next presented a prompt to enable Windows Update, where I clicked Yes. During the initial setup, the Windows firewall was set to deny all incoming connections. The setup process automatically resets the firewall settings to allow incoming connections after the first update.

This was the longest phase of the installation. This portion of the installation process installed and configured SQL Server and SharePoint and didn't require any additional user input. After this phase of the setup completed, the appliance rebooted and performed a software update.

The entire setup process took about an hour and a half. At that time, the BDA was completely joined to the domain, and SQL Server, SQL Server Analysis Services (SSAS), and SharePoint were all up and running. The installation process was amazingly simple.

If you've ever installed all these server products, you know that their installation processes are quite involved. The BDA took a surprisingly short time to install and configure the three server products. The appliance's installation process essentially set up Server 2008 R2 with the Web Server and Application Server roles installed. It also installed the Windows Server Backup feature. The BDA installation configured SharePoint as a single server web farm and installed SQL Server with a default instance name of POWERPIVOT.

Up and Running

After finishing the installation, you use the BDA's web-based SharePoint interface to work with the appliance. Users can use the appliance's network name to connect to it. For instance, in my tests, pointing the browser to http://BDA01 displayed the HP Business Decision Appliance home page. The home page is displayed by default and has links that let you download trial copies of Microsoft Office 2010, PowerPivot for Excel 2010, and Microsoft Silverlight. It also provides links to resources to help you learn more about PowerPivot, including a

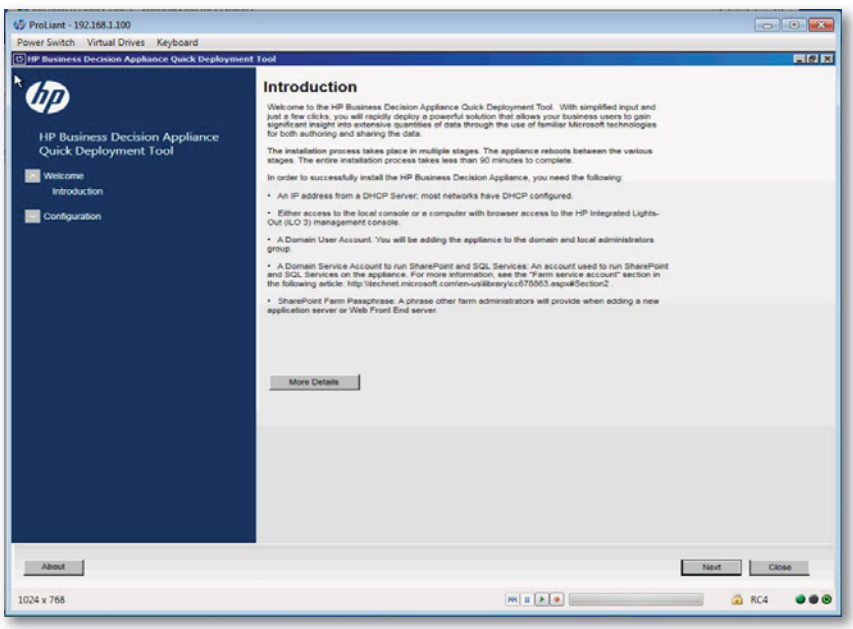

Figure 2: HP Business Decision Appliance Quick Deployment Tool

HP BUSINESS DECISION APPLIANCE

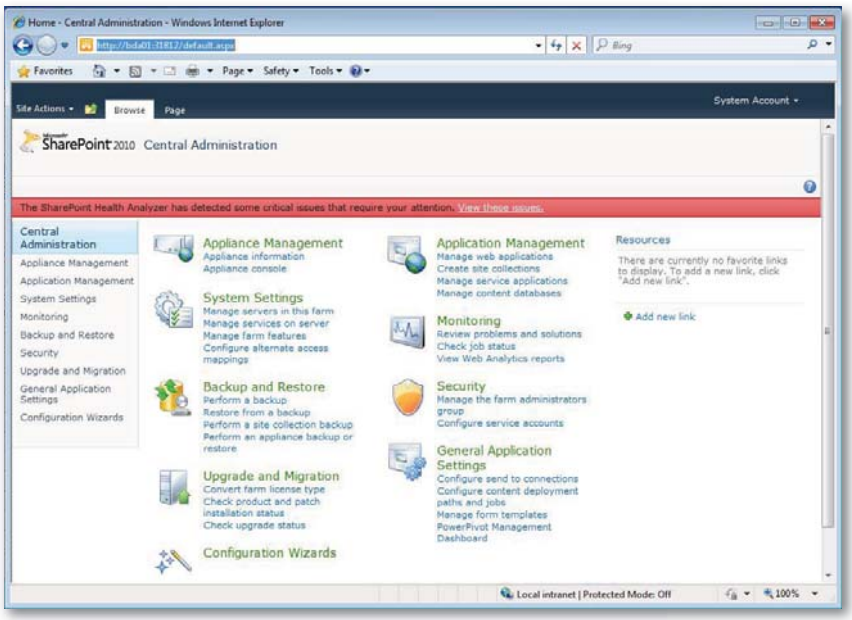

Figure 3: BDA SharePoint 2010 Central Administration page

virtual lab. You can later replace this default home page with one of your own if you desire.

Administrators can connect to the appliance by pointing their browsers to http://bda01:31812, which displays the SharePoint 2010 Central Administration page, as Figure 3 shows. The BDA Central Administration console is very much like the standard SharePoint 2010 Central Administration console except for the fact that it includes an Appliance Management link.

The Appliance Management link lets you view the appliance status and uptime, as well as shut down and back up the appliance. It also provides a link to perform a factory reset of the appliance, which restores everything—including the disk partitions—to the original state. You can use Microsoft System Center Operations Manager (SCOM) to manage the BDA. However, unlike a standard server, SCOM shows it as a new appliance icon. SCOM also checks the health of all three of the different server components.

IT professionals who create PowerPivot workbooks need to have Office 2010 installed; plus, they need to have PowerPivot for Excel. End users can consume the PowerPivot workbooks that are stored on the BDA simply by pointing their web browsers to the BDA and clicking on the links. They don't need to have Office 2010 or any other software in order to be able to open and use PowerPivot workbooks. Excel Services on the BDA takes care of all the rendering for browser-based client connections. The upshot is that you don't need to upgrade your entire client infrastructure to take advantage of the BI data the BDA provides.

To test the appliance, I created a multi million-row PowerPivot workbook and accessed it from the network. As you might expect, the appliance provided subsecond response times for this light load. For example, a 3.5-million row PowerPivot import took about 10 seconds. Subsequent work with the PowerPivot workbook was essentially instantaneous—including actions such as creating calculated columns on 3.5 million rows.

The BDA product managers I spoke with told me the device was designed to support 60 to 80 concurrent PowerPivot connections, which should equate to hundreds of real-world end users. Although I wasn't able to perform load testing on the appliance, Microsoft shared some internal LoadRunner performance test data that showed the appliance can run complex workloads from 65 unique connections using a variable 5- to 30-second think time, with response times in the 2-second range and about 14 percent CPU utilization. Heavy stress tests with a simpler workload supported as many as 250 users, with a response time of less than 3 seconds. The BDA undoubtedly delivers first-class performance.

Licensing

As you might expect, licensing for the appliance is a bit different from licensing a typical server. The base license for the BDA starts at $27,916. This price doesn't include the cost of SQL Server 2008 R2 Enterprise Edition or SharePoint 2010. This licensing structure allows businesses that already have volume agreements for SharePoint or SQL Server to apply those licenses toward the appliance.

Organizations that don't have volume license agreements can purchase SharePoint 2010 Enterprise and SQL Server 2008 R2 Enterprise licenses when they purchase the appliance. The BDA comes with a three-year service agreement, and HP is the single point of contact for all service requests.

Big Performance in a Tiny Package

The HP BDA provides incredible performance in a tiny 1U package. If your company has been looking to make the move to BI but has lacked the expertise to set up a scalable and secure BI infrastructure, you should definitely put the BDA at the top of your list. The appliance can be deployed in a couple of hours, and it provides first-rate performance.

InstantDoc ID 129569

HP Business Decision Appliance

PROS: Excellent performance in a small form factor; integrated management with SCOM; excellent out-of-band management with iLO; easy deployment; flexible licensing

CONS: Only one USB port; requires Office 2010 for content creation

RATING: ◆◆◆◆◆

PRICE: Starts at $27,916 for customers with SQL Server 2008 R2 Enterprise Edition and SharePoint 2010 Enterprise Edition

RECOMMENDATION: If your organization wants to take advantage of BI but doesn't have the expertise to configure the platform, or you're looking to deploy scalable BI solutions in your branch offices or departments, then you should definitely take a serious look at the HP BDA.

CONTACT: HP • 800-752-0900 • www.hp.com

Network Monitoring from Your Smartphone

Systems administration from the palm of your hand

by Eric B. Rux

If your job requires you to be on call, you probably carry a beeper or pager, or you have a phone that receives text messages when a mission-critical server or service goes down. But you really need to know what the entire network is doing to properly diagnose the problem. If you were at the office, you would fire up your network monitoring system and zero in on the problem. But if you're like me, these things usually happen when you're on the 14th hole of a disastrous golf game. It sure would be nice to know what the real problem was, and to know without a doubt if the problem were truly serious.

Several products let you monitor your network's status from your smartphone. Two of the products I reviewed, Paessler's iPRTG and GroundWork Open Source's Brooklyn For Nagios, work exclusively on Apple devices (iPod, iPhone, and iPad). The third product, ManageEngine's OpManager Smartphone GUI, can be used on most mobile browsers, including Apple products. I used an Apple iPhone 3GS to test each product. The two iPhone-ready apps point to a demo back-end monitoring system by default, so I could easily put the applications through the paces. But for the OpManager Smartphone browser-based application, I had to set up a complete ManageEngine OpManager server to test the application.

Installing the Apple-only products (iPRTG and Brooklyn For Nagios) involved using the Apple iTunes Store to redeem a coupon that you receive after purchasing the software. OpManager Smartphone GUI is accessed through your favorite mobile browser; you don't have to install any software on your phone.

Each product is written for a specific company's back-end monitoring solution. In other words, iPRTG is designed to work with PRTG Network Monitor, Brooklyn For Nagios is designed to work with Nagios, and OpManager Smartphone GUI is designed to work with OpManager. You can't use one company's iPhone app to hook up to another company's solution. Therefore, these mobile monitoring solutions should be part of your overall evaluation when you select a network monitoring solution.

All three products were very intuitive and easy to navigate. I never became lost in a maze of options or menus. Although there were clear differences between what the products could do, I found them to be very similar in overall function.

Each of the products I reviewed is designed to consume the content of back-end network monitoring solutions. You can't set up a new server or service to be monitored, nor can you add new users to the back-end monitoring system. Think of these tools as read-only.

Each product connects to its back-end monitoring system over the Internet, either on port 80 or port 443, and requires a username and password. Just as for any service that you use over the Internet, you need to take precautions to keep your company data secure, such as using SSL (port 443) and a strong username/password policy.

iPRTG

iPRTG uses a simple interface that packs in an abundance of information. For example, if you're monitoring a website, iPRTG shows you the website's address, the loading time of the website (in milliseconds), how long it has been up, when it was down last, and the overall uptime statistics.

As Figure 1 shows, the home page shows a list of favorites, which keeps your most important servers' information at your fingertips. If a server indicates a problem, iPRTG shows you exactly what the problem is and can even display the same graph as the full GUI product shows. In addition, iPRTG has live graphs and snapshots of data from 2 days, 30 days, and 1 year ago.

iPRTG supports multiple accounts, which you can save so that you don't have to log on to each one. For networks with a lot of systems, iPRTG includes a handy search feature that lets you zero in on the specific device you're looking for.

When you find you have a problem with a server, you can acknowledge the error on your iPhone. If you need to dive in further, the iPhone app has a link to the main PRTG web interface, which is similar to the desktop interface and allows you to perform advanced functions.

Figure 1: iPRTG home page

iPRTG

PROS: Favorites feature helps you monitor critical servers; includes graphs; can play a sound or vibrate on alarms

CONS: iPhone/iPod/iPad-only app (but does include a low-bandwidth browser for non-iPhones)

RATING: ◆◆◆◆◇

PRICE: $11.99 from iTunes

RECOMMENDATION: If you're already using PRTG Network Monitor, then iPRTG would be a great addition.

CONTACT: Paessler • 49 911 93775 0 • www.paessler.com

Brooklyn For Nagios

Brooklyn For Nagios has the simplest interface by far—mainly because it lacks the other two products' rich features. Instead of deep details, Brooklyn displays only whether the host is up or down and when it was last checked, as Figure 2 shows.

Just about the only feature Brooklyn has is the ability to acknowledge an outage. After you do this, it's time to break out the laptop, connect the phone for Internet access, and fire up the VPN. This is the only way you can dig deeper into the problem; Brooklyn doesn't have graphs or any other troubleshooting tools to help you narrow down the problem.

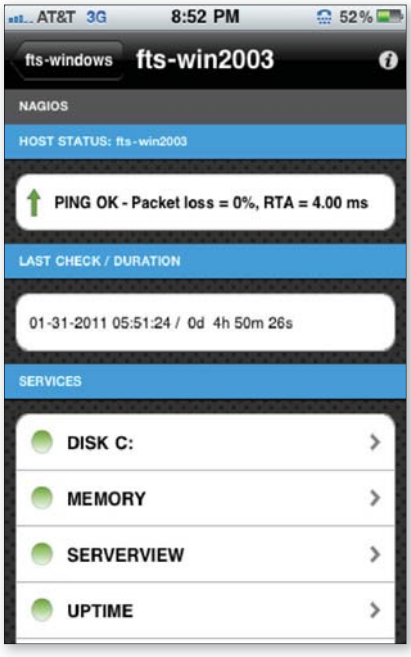

Figure 2: Brooklyn For Nagios interface

Does this mean Brooklyn For Nagios is a bad product? No, and here's why: Brooklyn's back end is Nagios, which is an open-source platform. This fact might make Brooklyn the ideal choice if you're looking for a mobile view of an inexpensive monitoring solution. Although Brooklyn itself isn't free, Nagios is.

Brooklyn For Nagios

PROS: Supports the freely available Nagios open-source network monitor software; extremely simple interface

CONS: iPhone/iPod/iPad-only app; not nearly as feature-rich as the other two products reviewed

RATING: ◆◆◆◇◇

PRICE: $12.99 from iTunes

RECOMMENDATION: If you're already using Nagios, then think seriously about adding this utility to your iPhone—just be sure to bring your laptop and a VPN connection with you on the golf course.

CONTACT: GroundWork Open Source • 866-899-4342 • www.groundworkopensource.com

OpManager Smartphone GUI

I was worried about reviewing OpManager Smartphone GUI. How could an HTML product compete with a native iPhone app? The truth is, it competes well—and in many cases, it's actually better.

The product's home page has six simple icons. Two of the icons, Down Devices and All Alarms, immediately indicate whether there's a problem. The list of devices includes indicators for the kind of device (e.g., Microsoft, Dell), as well as CPU and memory utilization.

When a device goes down, OpManager Smartphone GUI gives you two troubleshooting tools, Ping and Trace Route, that you can use to help narrow down the actual problem. These aren't last-result ping results, but on-the-fly pings that give you an immediate response.

OpManager Smartphone GUI lacks the cool graphs that the full product includes, instead replacing them with simple number indicators, as Figure 3 shows.

OpManager Smartphone GUI

PROS: Browser-based interface increases the chance it will work on any smartphone versus only on a specific platform; great home-page view

CONS: No favorites capability; lacks the utilization graphs that the full application boasts

RATING: ◆◆◆◆◆

PRICE: Included free with OpManager

RECOMMENDATION: If you're in the market for a network monitoring solution, give OpManager a close look; the product for mobile phones is outstanding.

CONTACT: ManageEngine • 877-386-3763 • www.manageengine.ca

Figure 3: OpManager Smartphone GUI interface

Editor's Choice

This review is less focused on which product is better than the others and more focused on simply letting you know what's available. If you're already using a particular product for network monitoring, your choices for an app are limited to the solutions that vendor offers. However, of the products I reviewed, OpManager Smartphone GUI gets the Editor's Choice award for smartphone-capable network monitoring.

InstantDoc ID 129645

Eric B. Rux
(ebrux@whshelp.com), Windows Home Server MVP, is a contributing editor for *Windows IT Pro* and writes a monthly column at svconline.com/connectedhome/windowshomeserver. Eric is the manager of technical support services at Eastern Washington University.

Third-Party Deployment Tools for Windows 7

Don't waste precious time suffering through manual OS deployments

by Blair Greenwood

The process of manually rolling out an OS across a network can be extremely time-consuming and complicated. Fortunately, if you're thinking about deploying Windows 7 across your organization's network, the manual route isn't your only option. There are many OS deployment tools available from third-party vendors that can help automate this tedious task, making it simpler and more convenient.

In addition to the ability to deploy Windows 7, these deployment tools provide an array of supplementary features that can assist you. Some of the features to consider include virtual machine (VM) deployment, management system integration, remote management, Active Directory (AD) integration, and Group Policy integration.

The key to finding the right solution is determining the perfect mix of features that your organization needs. To start your decision-making process, refer to the accompanying buyer's guide table, which outlines various Windows 7 deployment tools.

Key Considerations

Of course, Microsoft's System Center tools can greatly help with your deployment. It's likely that you already license some System Center products. However, System Center can be complex and its capabilities might be too limited for your needs. This is where third-party deployment tools are effective in providing solutions that are tailored to your needs. Deploying non-Windows OSs is another situation in which System Center might not provide all the functionality you desire, so be sure to determine whether System Center, or any third-party tool, will help deploy the OS that your organization requires.

If you're deploying only Microsoft OSs, Windows Deployment Service (WDS) might be a good choice for you. WDS is extremely useful because it lets you deploy Windows OSs remotely. Another useful WDS feature is its ability to set up clients over a network instead of installing an OS via CD-ROM or DVD. For more information about WDS installation and deployment, refer to Rhonda Layfield's article "Windows Deployment Service in Server 2008 R2" (December 2010, InstantDoc ID 125867).

Before choosing a deployment tool, ask yourself how your organization will be using virtualization. There's a good chance that you'll have to pay more for virtualization capabilities. However, it might be worth the additional cost, as virtualization is becoming an important and necessary tool in IT infrastructures. The ability to handle thin clients, application virtualization, and client virtualization can prepare you for the future by squeezing more out of your hardware now. Plenty of solutions that include virtualization support are available.

AD integration might be a worthy feature, providing all your deployment settings in a central database. AD is useful for both small and large operations, thanks to its ability to scale up or down easily. Another consideration is whether the deployment tool offers Group Policy integration. Group Policy management can be especially difficult to manually maintain.

If your organization is planning to upgrade from Windows XP to Windows 7, you have migration functionality to think about. There isn't an easy solution for this particular upgrade. One option is to reformat your hard drive and install Windows 7 on a clean new slate. However, you can save user data and settings through migration, making the process much easier. Many third-party solutions can migrate user data, application settings, and configuration settings. However, each deployment tool is unique and might offer only a fraction of the migration capabilities that your organization requires. Therefore, it's critical to determine whether your needs correspond with the deployment tool's migration offerings.

The Right Mix

Choosing the right deployment tool often depends on the size of your business. A small business might be able to manually deploy its OSs with little hassle. However, if you're looking at a large enterprise deployment, a third-party deployment tool could be worthwhile. Although these third-party deployment tools might be a more expensive prospect than deploying your OSs manually, it might be worth those extra dollars for the convenience of an automated rollout. Refer to the buyer's guide table for a comparison of various Windows 7 OS deployment tools that will help you decide which one is right for you.

InstantDoc ID 129668

BLAIR GREENWOOD
(blairg@rams.colostate.edu) is a contributing editorial intern for *Windows IT Pro*. She is a senior technical journalism major at Colorado State University. She's worked with Java, C, and Visual Basic programming languages.

Company	Product	Pricing	OS Deployment?	Application Deployment?	VM Deployment?	Management System Integration?
Acronis 781-782-9000 877-669-9749 www.acronis.com	Acronis Backup & Recovery 10	$74 for workstations and $853 for servers	Windows (7, XP, Server 2008/2003, 2000), VMware ESX, Linux	Yes	Yes	Yes
	Acronis Snap Deploy 3	$25 for PCs, $121 for servers	Windows (7, XP, Server 2008/2003, 2000), Small Business Server 2003, Linux	Yes	Yes	Yes
	Acronis True Image Home 2011	$49.99	Windows (7, Vista, XP)	Yes	Yes	No
	Acronis True Image Home 2011 Plus	$79.99	Windows (7, Vista, XP)	Yes	Yes	No
CA Technologies 800-225-5224 www.ca.com	CA Client Automation	$85 per managed system with tiered volume discounts	Windows (7, CE, Mobile), Mac OS X, Linux	Yes	Yes	Yes
	CA Plex	$6,000 per developer	Windows 7, J2EE/Java, IBM I/AS400	Yes	Yes	No
Dell KACE 877-646-8366 www.kace.com	Dell KACE System Management Appliances	Starts at $4,500	Windows 7, Mac OS X, Linux (Red Hat 3, 4, 5)	Yes	Yes	Yes
HP 800-289-6947 www.hp.com	HP Client Automation Enterprise Edition	Between $54 and $109 per managed endpoint	Windows (7, Vista, XP, 2000), Linux (SUSE, Red Hat)	Yes	Yes	Yes
Novell 801-861-4272 800-529-3400 www.novell.com	Novell ZENworks Configuration Management 11	$85 per seat	All Windows and Linux desktop OSs	Yes	Yes	Yes
Prowess 206-442-1117 888-733-7569 www.smartdeploy.com	SmartDeploy Enterprise	$2,294.95	Windows (7, XP, 2000, Server 2008/2003)	Yes	Yes	Yes
ScriptLogic 561-886-2400 800-813-6415 www.scriptlogic.com	Desktop Authority	$39 per seat	Windows (7, Vista, XP, Server 2008/2003, 2000)	Yes	No	No
Symantec 650-527-8000 800-745-6054 www.symantec.com	Altiris Client Mangement Suite 7	$95 per node	Windows (7, XP, Server 2008/2003/2000), Linux (Red Hat, SUSE), VMware ESX, Solaris	Yes	Yes	Yes

Editor's Note: Information in this buyer's guide comes from vendor representatives and resources and is meant to jumpstart, not replace, your own research; also, the table isn't necessarily comprehensive, as some products might have been left out due to the writer's oversight or a lack of vendor response.

	Migrate User Data?	Migrate Application Settings?	Migrate Configuation Settings?	Remote Management?	AD Integration?	Group Policy Integration?	Types of Deployment Images	Ability to Deploy .NET Framework?	Ability to Deploy Patches to Applications?	Ability to Deploy Patches to OSs?
	Yes	Yes	Yes	Yes	Yes	Yes	.tib	No	Yes	Yes
	Yes	Yes	Yes	Yes	Yes	No	.tib	No	Yes	Yes
	Yes	Yes	Yes	No	No	No	.tib	No	Yes	Yes
	Yes	Yes	Yes	No	No	No	.tib, .zip	No	Yes	Yes
	Yes	Yes	Yes	Yes	Yes	Yes	.wim, .gho	Yes	Yes	Yes
	No	No	No	Yes	No	No	n/a	Yes	Yes	Yes
	Yes	Yes	Yes	Yes	Yes	No	.wim, WIN images, network scripted installs, and customer K-images	Yes	Yes	Yes
	Yes	Yes	Yes	Yes	Yes	Yes	.wim	Yes	Yes	Yes
	Yes	Yes	Yes	Yes	Yes	Yes	WinPE, Ghost, ZENworks, ImageX	Yes	No	No
	Yes	Yes	Yes	Yes	No	No	.wim	Yes	No	No
	Yes	Yes	Yes	Yes	Yes	Yes	n/a	Yes	Yes	Yes
	Yes	Yes	Yes	Yes	Yes	Yes	.wim, .gho, .img, .Solaris	Yes	Yes	Yes

INSIGHTS FROM THE INDUSTRY

Free Tool: The Performance Analysis of Logs (PAL) Tool

Frustrated by the time-consuming task of analyzing Performance Monitor data? It usually means you have to study Performance Monitor logs in graph view, and hope that you have the expertise to understand the voluminous information in front of you. In these days of dwindling IT resources, it's a tough proposition.

Now, more than ever, a tool such as Performance Analysis of Logs (PAL) is indispensable. Free from Microsoft's CodePlex site, PAL analyzes your performance counter logs and provides HTML threshold reports. You can use preset threshold counters (available for most major Microsoft products) or customize your own counters.

We've covered PAL in the magazine before, and it's worth looking at the tool's previous coverage. In "Get a Handle on Windows Performance Analysis" (www.windowsitpro.com, InstantDoc ID 101162), Michael Morales writes the following:

"A tool that Microsoft support relies on to analyze Performance Monitor logs is the Performance Analysis of Logs (PAL) Tool. Clint Huffman, a Microsoft senior premier field engineer, wrote the 6,000-line VBScript tool, which is free and open source. PAL lets administrators easily analyze Performance Monitor logs without requiring them to be experts in performance counters or Windows architecture.

"PAL contains a wizard-based UI that asks specific information about the system, which PAL passes as arguments to the VBScript program. PAL picks up where other log analyzers leave off, such as taking into account whether the system is 64-bit or 32-bit, whether the /3GB switch is used, and how much physical memory is installed—all variables that affect system performance. PAL uses these variables along with known thresholds, which were determined by engineers with years of experience, to determine the analysis that's displayed. PAL provides a chronological order of alerts, so that you can correlate your system's performance to any problems that you noticed at specific times.

"PAL also can provide application-specific analysis for applications such as Microsoft BizTalk Server, Microsoft Exchange Server, Microsoft Office SharePoint Server, Microsoft SQL Server, and Microsoft IIS. So as an administrator wearing several hats, you can have application-specific performance data analyzed without being an expert in the performance counters for an application. PAL can make your life easier by providing analysis for baseline data when performance is typical or to help pinpoint the root cause of a performance issue when a problem occurs.

PAL analyzes your performance counter logs and provides HTML threshold reports. You can use preset threshold counters or customize your own counters.

"PAL's user-friendly UI walks you through the few steps necessary to start the analysis process. The analysis report that PAL generates is an .html file that's stored by default under the My Documents\PAL Reports folder. The report contains hyperlinks and graphs that enable easy interpretation and navigation, and the file's portability lets you easily store it in a convenient location."

In "Two Exchange Server Tools You Should Know About" (InstantDoc ID 100132), Paul Robichaux focuses on the tool's Exchange benefits:

"[PAL] is a free tool available from Microsoft's CodePlex. The concept behind PAL is simple: It ingests Performance Monitor log files from a server running Exchange, Microsoft IIS, Microsoft BizTalk Server, or several other Microsoft applications, then provides charts, graphs, and alerts for the most significant application-related performance counters. PAL is driven by XML files that specify which counters are important for a particular application. The tool ships with an XML file for Exchange 2003 that's based on Microsoft's 'Troubleshooting Microsoft Exchange Server Performance' white paper. The Exchange 2007 counter file is based on the list of counters in the Exchange documentation article 'Monitoring Without System Center Operations Manager.'

"PAL is straightforward to use: You install it (and its prerequisites, which include Log Parser, the .NET Framework, and the Office 2003 web components), then run it and use its wizard-like interface to select the counter file you want to process and, if you're on Exchange 2007, the server role that generated it. Depending on the role you choose, you might need to answer questions about the selected server, such as the number of processors or the amount of RAM. After you've done so, PAL processes the logs and renders HTML reports that highlight the most notable performance data from the server. You can use PAL to get a quick overview of how an Exchange Server is performing, then drill down to get more detail on specific counter sets that give information about a specific subsystem or component."

Download PAL today at pal.codeplex.com.

—Jason Bovberg

INDUSTRY BYTES

Push Android Apps to Smartphones and Tablets

In an enterprise setting, users have become accustomed to being handed a laptop that is configured, has all the permissions and access levels it needs, and all the software necessary for day-to-day work. Why should a smartphone or tablet be any different?

Zenprise has introduced something fairly revolutionary in this regard with the Android Enterprise Application Store. The concept is what Zenprise calls "the first Android enterprise app store." What it basically means is that IT will be able to push relevant applications to users so the minute they start up a company smartphone or tablet, the programs they need are right at their fingertips.

Additionally, the app store lets you set up groups—for example, you could configure it so that any tablet that is configured as part of the "marketing group" will automatically get a certain set of apps. You can do the same thing with documents, meaning you can easily push a PowerPoint slide to the 25 individuals in sales with a few clicks, for example. (You can also import your Active Directory groups into Zenprise for quick classification.)

This whole vision assumes to some extent that your company is "on board" with the idea that smartphones and tablets (sponsored by the company) are the wave of the future. It makes sense if either (1) the majority of your organization is equipped with mobile devices or (2) at least one or more departments are equipped with said devices.

Zenprise has also announced remote control for Android, which essentially gives the same remote control capabilities the company has already offered for BlackBerry and Windows Mobile.

Key new features of the solution include:

- **Task manager.** Zenprise introduces a task manager very similar to that in Windows, so you can quickly troubleshoot an Android user's problem and see if an app has stalled (or if they just had 10 apps running in the background, which would explain slow performance).
- **Chat.** IM from your PC with an Android user in case a phone call isn't realistic (e.g., the user is in a meeting).
- **Remote reboot** for a user's device.
- **Remote control for Android devices** to troubleshoot problems.
- **Push documents to the device** on an individual basis or based on groups.
- **Users don't have to go to the marketplace**, at least for applications required for work. Users might still want games or other fun apps for home use, but you could choose to block network access and only install apps based on request, if you're concerned about security and the many sketchy apps on the Android Marketplace.

The features I have outlined have been rolled up as part of the existing Zenprise product and are compatible with all versions of Android 1.6 and up. To learn more about pricing or other details, visit www.zenprise.com.

—Brian Reinholz

Wary of Cloud Backup? How About "Private Cloud" Backup?

According to recent surveys from leading storage and backup vendors, one of the primary concerns among IT admins regarding cloud backup is security. The fear may be unfounded, however: According to the same surveys, once a given company actually engages the services of a cloud backup provider, and lives with the new reality for a while, the fear subsides dramatically. Still, the "fear of the unknown" can be powerful, and many businesses remain wary of handing over their data to a third party.

As if responding to such concerns, 3X Systems has introduced its 3X RBA Enterprise Series, a unique remote backup appliance. The company calls the technology behind the appliance a "private cloud" architecture, which ensures that users' data is safe and easy to recover, yet confidential, and always under their own control. According to 3X Systems, these are "critical considerations for organizations such as those in health care, financial services, non-profit, and professional services that need to comply with regulations governing electronic data."

With 3X RBA Enterprise Series, larger distributed organizations can back up all their Windows-based servers, workstations, and laptops over the Internet to a central, easy-to-manage appliance that delivers complete data protection and disaster-recovery capabilities. Integrated de-duplication, encryption, and block-level backup capabilities mean that data changes are securely and efficiently transmitted to the central 3X RBA appliance. The Enterprise Series provides up to 10TB of usable storage. The system offers the following features:

"Bare Metal Recovery"—By protecting the entire system image, including all settings, applications, and data, user productivity is restored much faster after suffering data loss, theft, or corruption. In contrast, data-only recoveries require users to spend hours, if not days, reinstalling software applications and reconfiguring system settings that ultimately may never be quite the same.

Support for latest Exchange Server backup and "brick-level" recovery—Administrators and users can now quickly and easily recover granular data ("bricks"), such as an individual mailbox, email message, calendar, or contact, without having to recover the entire Exchange system.

Support for IT environments where computer resources are shared using virtualization. Faster data transfer rates—3X Systems has accelerated the data transfer rate between the protected computer and the remote appliance for faster backup and recovery.

For more information about the "private cloud," visit www.3x.com.

—Jason Bovberg

1&1® WEB HOSTING

PROFESSIONAL WEBSITES

As the world's largest web host, we know the developer features you need in a hosting package!

.com
.info .org
.net

Domains Included
All hosting packages include domains, free for the life of your package.

Unlimited Traffic
Unlimited traffic to all websites in your 1&1 hosting package.

Developer Features
Extensive language support with PHP 5/6 (beta) with Zend Framework and git version management software.

Online Marketing Tools
SEO tools to optimize your website.
1&1 Webstatistics makes it easy to monitor your progress.

NEW! Now offering .ca domains to our Canadian customers. C$9.99/first year.* Visit www.1and1.ca for details.

1&1® HOSTING PACKAGES
6 MONTHS FREE!*
OFFER EXTENDED!

1&1® BUSINESS PACKAGE:

- 3 Included Domains
- Private Domain Registration
- 250 GB Web Space
- UNLIMITED Traffic
- NEW: Version Management Software (git)
- 2,500 E-mail Accounts
- 50 MySQL Database
- 25 FTP Accounts
- E-mail Marketing Tool
- 24/7 Toll-free Customer Support

~~$9.99 per month*~~ ~~C$9.99 per month*~~

Need more domains?
.info domain only $0.99/first year*
.com domain only $4.99/first year*

More special offers available on our website!

MEMBER OF united internet

1-877-GO-1AND1 www.1and1.com
1-855-CA-1AND1 www.1and1.ca

1&1®

*Offers valid for a limited time only. 12 month minimum contract term applies for web hosting offers. Setup fee and other terms and conditions may apply. Domain offers valid first year only. After first year, standard pricing applies. Visit our website for full promotional offer details. Program and pricing specifications and availability subject to change without notice. 1&1 and the 1&1 logo are trademarks of 1&1 Internet AG, all other trademarks are the property of their respective owners.

Bringing OWA to Your Desktop

When Microsoft Exchange Server 2010 launched over a year ago now, one of its selling points was the vast improvements in Outlook Web App, bringing the web client on par with the desktop version of Microsoft Office Outlook. In addition to getting the rich-client experience on multiple browsers (Firefox and Safari being added to Internet Explorer), OWA 2010 gave users their first look at Exchange 2010 features such as Conversation View and MailTips, which otherwise you need Outlook 2010 to access.

At the time, I wrote about the possibility that the improved nature of OWA would lead companies to adopt OWA as their standard for email and avoid upgrading to the latest desktop Outlook—and remember, Outlook 2010 wasn't available until about six months after Exchange 2010, giving early adopters of Exchange 2010 plenty of time to see if this could be a workable solution. I spoke with Messageware President and CEO Mark Rotman, and although he said he doesn't have hard data, he believes a significant number of companies are deploying OWA alone, at least to a segment of their users.

One of the nicest features is the full control you can get from the context menu in the Taskbar, including the ability to compose a message or schedule an appointment without going into the client window itself.

"With Exchange 2010 and enterprise class customers—5,000 or more employees—where OWA is deployed as an external connectivity mechanism for them," Rotman said, "I think every single one of those companies has at least one department or one group of users that will be OWA-only, and in many cases thousands of users that will be OWA-only." This situation works great for mobile or remote workers, but can be just as effective for regular office workers, and can save companies big money over deploying and managing desktop Outlook.

Messageware obviously knows a few things about OWA—they've been making products that both enhance security and add functionality to OWA for years. So perhaps it seems natural that, seeing this trend of companies deploying OWA without Outlook, the company would develop its own desktop client. The result, which at this point is still in beta, is OWA Desktop, and it's designed specifically to interoperate with users already on OWA but to provide a desktop experience—and do so at a price much reduced from purchasing the full Microsoft client.

Rotman walked me through a demo of OWA Desktop, and I have to admit it's pretty impressive. It gives you all the functionality you've come to expect and depend on from Outlook—calendar reminders, new mail pop-ups, and so forth. It also includes features that OWA and even desktop Outlook itself don't have. One of the nicest features is the full control you can get from the context menu in the Taskbar, including the ability to compose a message or schedule an appointment without going into the client window itself.

OWA Desktop lets an individual manage multiple email accounts. Importing and exporting contacts is also a simple procedure—and not something so easily done just with OWA. And it will work with on-premises Exchange servers or hosted versions of Exchange, whether from Microsoft's offerings or any third-party hoster—essentially, any version of OWA. As Rotman said, "It's all about convenience, all about working quickly, and all about having that full desktop-like functionality."

Messageware is planning a Q1 release for OWA desktop; when I spoke with Rotman, final pricing had not yet been set. An additional benefit over Outlook is a very small install size—less than 5MB. Visit www.messageware.com if you're interested in checking out the beta or any of the company's other OWA add-ons. And in the meantime, what are you doing to provide your users with the best email and calendaring features? Outlook 2010? Outlook 2007? Just OWA? Or something else altogether? Email bwinstead@windowsitpro.com and let me know.

—B.K. Winstead

AD INDEX

For detailed information about products in this issue of *Windows IT Pro*, visit the web sites listed below.

VENDOR DIRECTORY

The following vendors or their products are mentioned in this issue of *Windows IT Pro* on the pages listed below.

DIRECTORY OF SERVICES | WINDOWS IT PRO NETWORK

Search our network of sites dedicated to hands-on technical information for IT professionals.
www.windowsitpro.com

Support

Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.
www.windowsitpro.com/go/forums

News

Check out the current news and information about Microsoft Windows technologies.
www.windowsitpro.com/go/news

EMAIL NEWSLETTERS

Get free news, commentary, and tips delivered automatically to your desktop.

- *asp.netNOW*
- *DevProConnections UPDATE*
- *Exchange & Outlook UPDATE*
- *Security UPDATE*
- *SharepointPro Connections UPDATE*
- *SQL Server Magazine UPDATE*
- *Windows IT Pro UPDATE*
- *Windows Tips & Tricks UPDATE*
- *WinInfo Daily UPDATE*

www.windowsitpro.com/email

RELATED PRODUCTS

Custom Reprint Services

Order reprints of *Windows IT Pro* articles. Diane Madzelonka at Diane.madzelonka@penton.com.

Windows IT Pro VIP

Get exclusive access to over 40,000 articles and solutions on CD and via the Web. Includes FREE access to eBooks and archived eLearning events, plus a subscription to either Windows IT Pro or SQL Server Magazine.
www.windowsitpro.com/go/vipsub

SQL SERVER MAGAZINE

Explore the hottest new features of SQL Server, and discover practical tips and tools.
www.sqlmag.com

ASSOCIATED WEBSITES

DevProConnections

Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at DevProConnections.com, where IT pros creatively and proactively drive business value through technology.
www.devproconnections.com

SharePointPro Connections

Dive into Microsoft SharePoint content offered in specialized articles, member forums, expert tips, and Web seminars mentored by a community of peers and professionals.
www.sharepointproconnections.com

NEW WAYS TO REACH *WINDOWS IT PRO* EDITORS:

LinkedIn: To check out the *Windows IT Pro* group on LinkedIn, sign in on the LinkedIn homepage (www.linkedin.com), select the Search Groups option from the pull-down menu, and use "Windows IT Pro" as your search term.

Facebook: We've created a page on Facebook for *Windows IT Pro*, which you can access at: http://tinyurl.com/d5bquf. Visit our Facebook page to read the latest reader comments, see links to our latest web content, browse our classic cover gallery, and participate in our Facebook discussion board.

Twitter: Visit the *Windows IT Pro* Twitter page at www.twitter.com/windowsitpro.

Windows IT Pro

by Jason Bovberg

April Fools!

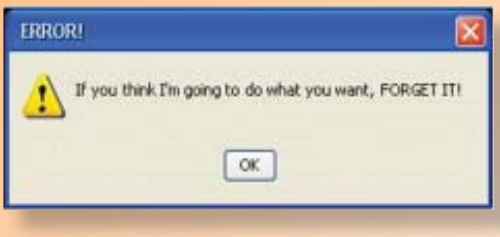

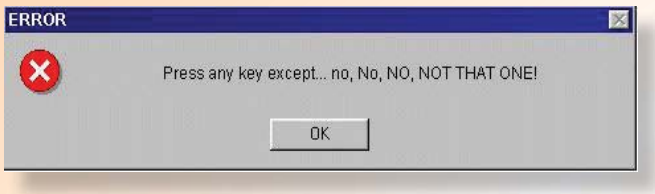

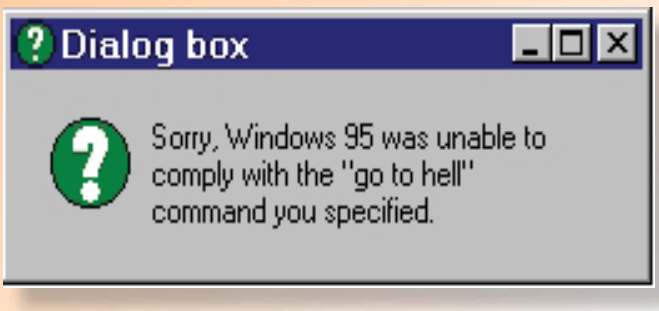

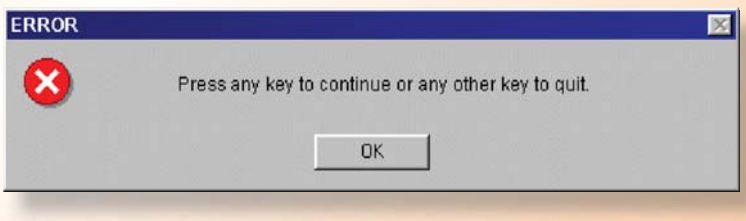

PRODUCT OF THE MONTH

Emoticons have become an essential part of our collective online conversation. But we all know how incredibly difficult it can be to type out each individual character that makes up that emoticon, right? At the Consumer Electronics Show in Las Vegas earlier this year, Russell Holly (Geek.com) caught sight of a product that will make all our communications easier. It's the USB Emoticon Keypad from Lavatelli. This photo shows a prototype model, but once this baby launches into production, we can all look forward to expressing the breadth of our emotional landscape with the tap of a button! Watch for it at www.bajca.com.

April 2011 issue no. 200, *Windows IT Pro* (ISSN 1552-3136) is published monthly. *Windows IT Pro,* 748 Whalers Way, Fort Collins, CO 80525, (800) 793-5697 or (970) 663-4700. Sales and Marketing Offices: 748 Whalers Way, Fort Collins, CO 80525. Advertising rates furnished upon request. Periodicals Class postage paid at Fort Collins, Colorado, and additional mailing offices. **POSTMASTER: Send address changes to *Windows IT Pro,*** 748 Whalers Way, Fort Collins, CO 80525**. SUBSCRIBERS:** Send all inquiries, payments, and address changes to *Windows IT Pro,* Circulation Department, 748 Whalers Way, Fort Collins, CO 80525. Printed in the USA.

SIMPLIFY

command of the data center with the power of convergence.

HP Converged Infrastructure simplifies what's next in management with HP BladeSystem.

Complexity drains your resources, leaving less for innovation.

Get it under control with the HP BladeSystem integrated, centralized management solution—HP Insight Control. With HP BladeSystem c7000 featuring the HP ProLiant BL460c G7 server powered by the Intel® Xeon® processor 5600 series, HP Insight Control can put you back in charge of your data center operations.

HP Insight Control delivers:

- Up to 40% increase in the productivity of system administrators*
- Up to 83% reduction in unplanned downtime*
- Payback in less than 5 months*

See how to take back control with the IDC white paper *Gaining Business Value and ROI with HP Insight Control Management Software.*
>> hp.com/go/takecontrol4

HP ProLiant BL460c G7 server

- Two six-core Intel® Xeon® processor 5600 series (2.53GHz) installed
- 16GB of memory; expandable to 384GB
- HP Smart Array P410i Controller
- One integrated NC553i Dual Port 10Gb FlexFabric Converged Network Adapter
- Up to two HP hot plug small form factor SAS, SATA, or Solid State drives

$4,759 (Save $467)

Lease for just $116 /mo.†

SmartBuy (PN: 630442-S01)

*Source: IDC white paper sponsored by HP *Gaining Business Value and ROI with HP Insight Control Management Software*, #224704, September 2010

†Prices shown are HP Direct prices; reseller and retail prices may vary. Prices shown are subject to change and do not include applicable state and local taxes or shipping to recipient's address. Offers cannot be combined with any other offer or discount and are good while supplies last. All featured available offers in U.S. only. Savings based on HP published list price of configure-to-order equivalent (HP ProLiant BL460c G7 server: $5,226 -$467 = SmartBuy price $4,759). Financing available through Hewlett-Packard Financial Services Company and its subsidiaries (HPFSC) to qualified commercial customers in the U.S. and is subject to credit approval and execution of standard HPFSC documentation. Prices shown are based on a lease 48 months in term with a fair market value purchase option at the end of the term and are valid through December 31, 2011. Other charges and restrictions may apply. HPFSC reserves the right to change or cancel this program at any time without notice.

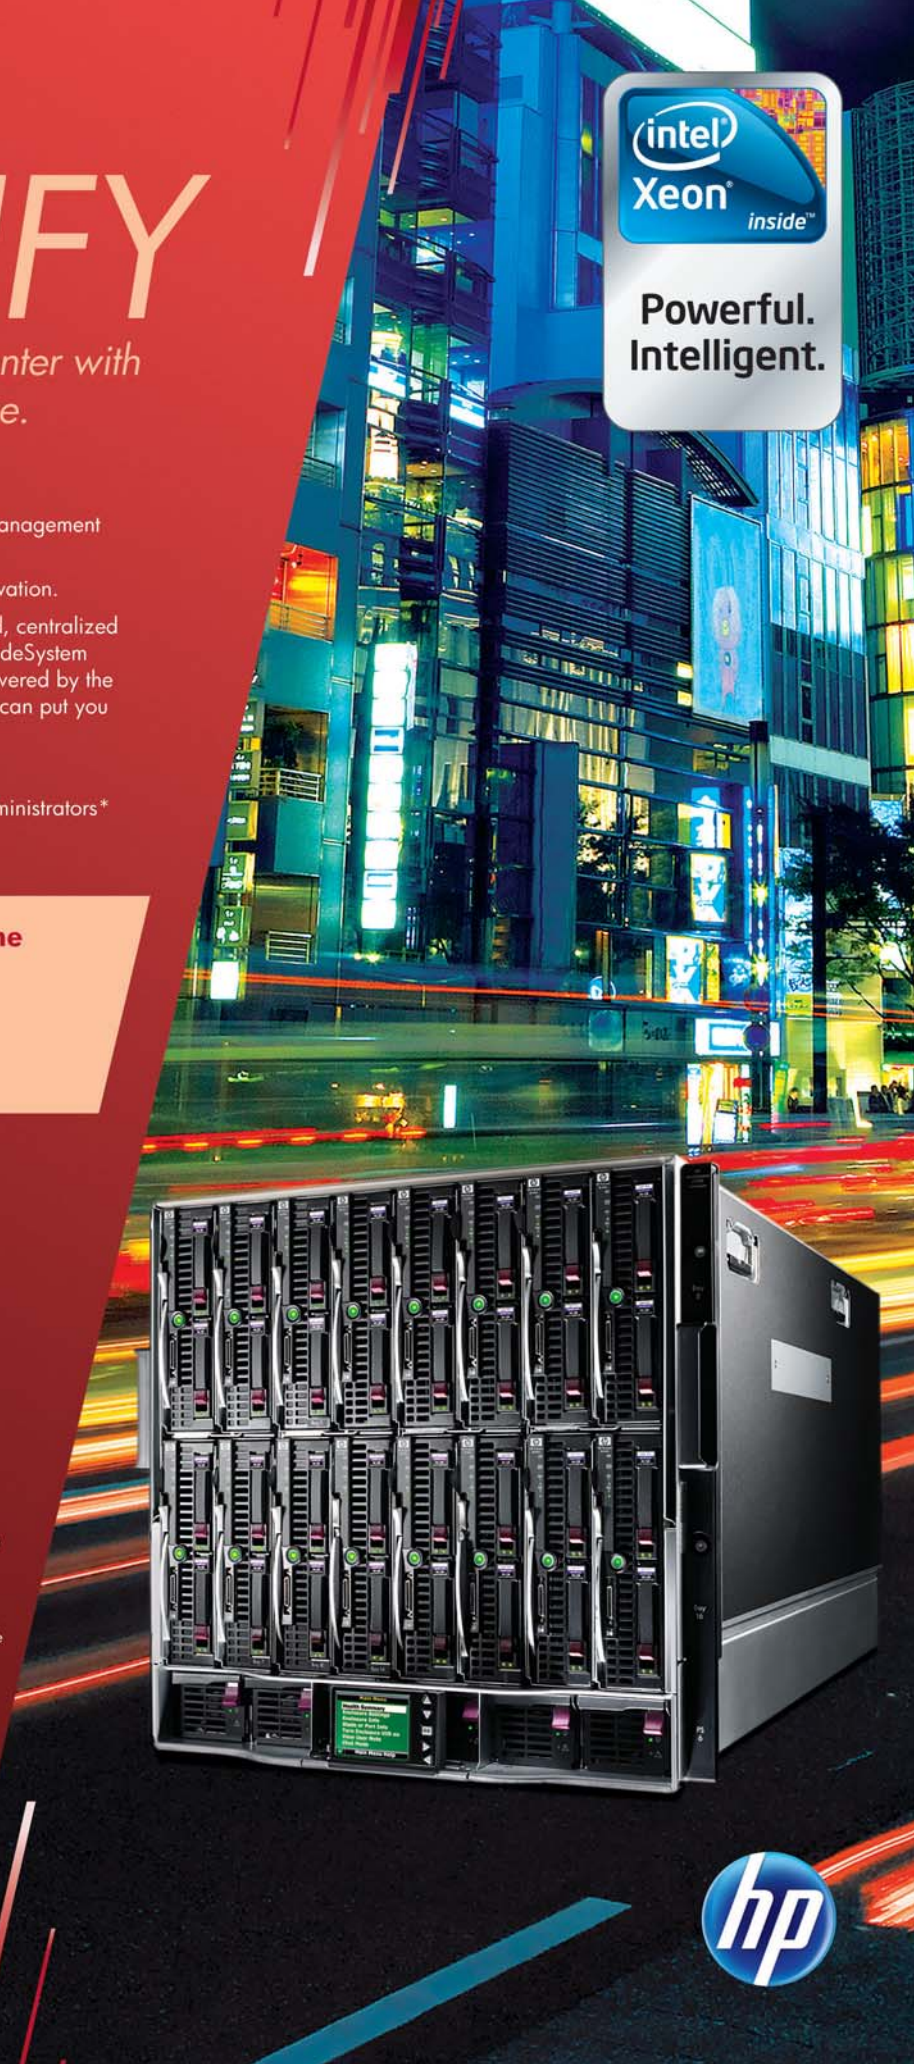

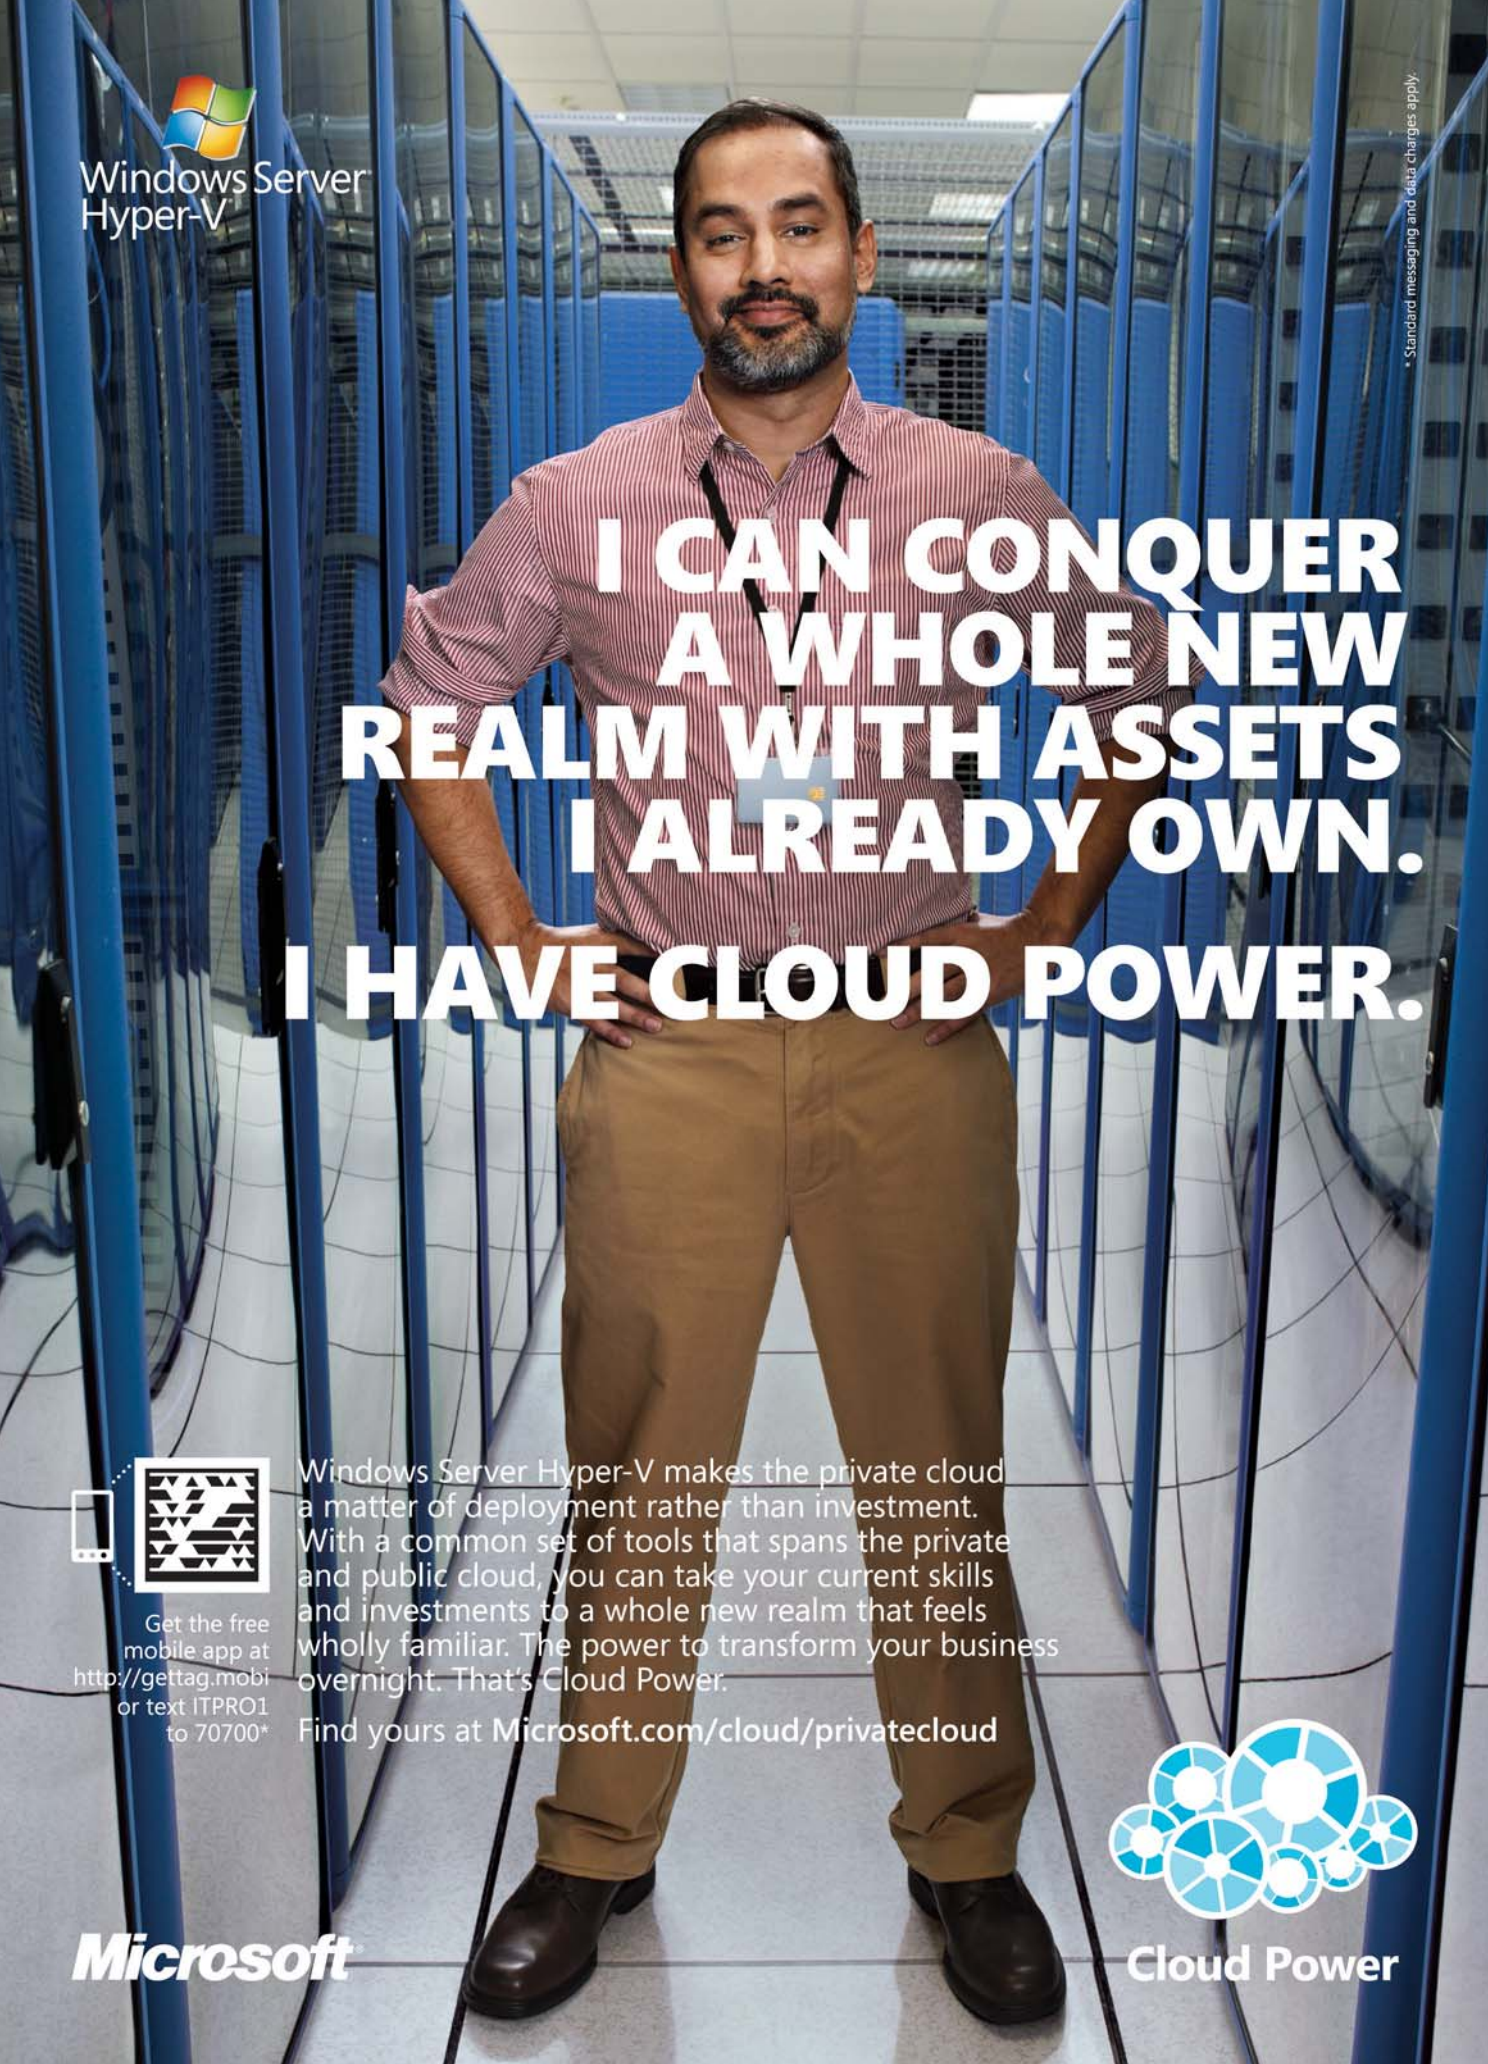
Windows Server Hyper-V
I CAN CONQUER A WHOLE NEW REALM WITH ASSETS I ALREADY OWN.
I HAVE CLOUD POWER.
* Standard messaging and data charges apply.
Get the free mobile app at http://gettag.mobi or text ITPRO1 to 70700*
Windows Server Hyper-V makes the private cloud a matter of deployment rather than investment. With a common set of tools that spans the private and public cloud, you can take your current skills and investments to a whole new realm that feels wholly familiar. The power to transform your business overnight. That's Cloud Power.
Find yours at Microsoft.com/cloud/privatecloud
Microsoft
Cloud Power